arXiv:1408.5544v2 [stat.ML] 27 Aug 2014

To lie or not to lie in a subspace

Daniel L. Pimentel-Alarcón
Electrical and Computer Engineering, Mathematics
pimentelalar@wisc.edu

RESEARCH PROPOSAL

University of Wisconsin-Madison
Madison, WI, 53706, USA

August 28, 2014

Submitted for revision to committee members:

Robert D. Nowak
Electrical and Computer Engineering
nowak@ece.wisc.edu

Nigel Boston
Electrical and Computer Engineering, Mathematics
boston@math.wisc.edu

Abstract

Give deterministic necessary and sufficient conditions to guarantee that if a subspace *fits* certain **partially observed** data from a union of subspaces, it is because such data really *lies* in a subspace.

Furthermore, give deterministic necessary and sufficient conditions to guarantee that if a subspace fits certain partially observed data, such subspace is *unique*.

Do this by characterizing when and only when a set of incomplete vectors *behaves* as a single but complete one.

Contents

1 Prologue **3**

2 Introduction **5**
2.1 Insight 6
2.2 The essence 6
2.3 Organization of the document 7

3 Preamble **8**

4 Setup **11**
4.1 Subspaces and bases 11
4.2 Vectors and bases 11
4.3 Observation sets 12
4.4 Incomplete vectors, bases and subspaces 13
4.5 Fitting incomplete vectors 15
4.6 Fitting generic vectors $\sim$ fitting observation sets 15

5 Assumptions **18**

6 Results **20**

7 Uniqueness **23**
7.1 The subspace $\mathcal{U}_\Omega$ 23
7.2 One at a time 24
7.3 Several at Once 25
7.4 Independence 27
7.5 All you need is $d - r$ 32
7.6 Proof of Theorem 2 33

8 All of a kind **34**
8.1 Characterization of $\mathcal{K}$ 36
8.2 All you need is $d - r$ of a kind 37
8.3 Bases Characterization 38
8.4 Proof of Theorem 1 38

9 Intuitively speaking **40**
9.1 Using Lemma 6 43
9.2 The idea behind Lemma 6 45
9.3 About our assumptions 48
9.4 Almost every $\neq$ every 51

10 Epilogue **54**

11 List of symbols, references and index **56**

1 Prologue

We love subspaces. We observe a phenomenon and try to find a line that explains it. We get our hands on some data, and we try to find a subspace that fits it. But what if we are looking for subspaces where there really are not? How can we guarantee that if we find a subspace, it is because there really *is* a subspace? In other words, how can we make sure that if certain data *fit* in a subspace, it is because it really *lies* in such subspace?

In many cases we don't really have to worry about this problem. For instance, if we have a collection of generic vectors that fit in an r-dimensional subspace, as long as our collection has more than r vectors, we can always verify if our collection indeed lies in an r-dimensional subspace, because we will always have an *extra*, *generic* vector to validate this. This is because almost surely, a set of more than r generic vectors fits in an r-dimensional subspace iff it actually lies in such subspace.

Nevertheless, if we suppose that our collection of vectors is only **partially observed**, this becomes a much harder problem, as a set of arbitrarily many *incomplete vectors* may fit in an r-dimensional subspace even if their *complete* counterparts do not really lie in a subspace.

Example 1. Suppose $r = 1$, and consider the following set of vectors:

$$\mathcal{X} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{bmatrix}.$$

It is easy to see that they do not lie in a 1-dimensional subspace. Nevertheless, suppose that we only observe a subset of their entries:

$$\hat{\mathcal{X}} = \begin{bmatrix} 1 & 1 \\ 1 & \cdot \\ \cdot & 3 \end{bmatrix}.$$

Then both *incomplete vectors* fit in the 1-dimensional subspace spanned by

$$U = \begin{bmatrix} 1 \\ 1 \\ 3 \end{bmatrix},$$

despite their full counterparts do not lie in a 1-dimensional subspace. ■

Of course, in general, without knowing anything a priori about our data there is no hope to succeed at this task, as the missing entries could be arbitrary. Fortunately there are ma cases of data that lies in —or can be accurately approximated by— a union of subspaces[1], a beautiful setup under which this task is not only feasible but also non-trivial. This is precisely the assumption under which we will operate, i.e.,

We will assume in the rest of this document that every vector of our data lies in the union of $\mathcal{S}^\star$, a set of r-dimensional subspaces of $\mathbb{R}^d$.

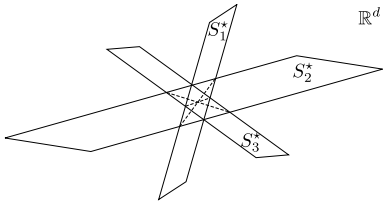

Figure 1: Example of the union of three 2-dimensional subspaces of $\mathbb{R}^3$.

At first glance this might deceivingly appear as a trivial task: if all vectors lying in a union of subspaces fit in one r-dimensional subspace, how could they not all lie in an r-dimensional subspace? To see this, consider the following.

Example 2. With the same setup as in Example 1, further suppose that $\mathcal{S}^\star$ is the set of the two subspaces spanned by the following vectors:

$$U_1^\star = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \qquad U_2^\star = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}.$$

If we again assume that our data is the set of vectors $\mathcal{X}$ —which clearly lies in the union of the subspaces of $\mathcal{S}^\star$— but that we only observe the subset of their entries in $\hat{\mathcal{X}}$, it is easy to see that U fits our data despite their full counterparts do not lie in a 1-dimensional subspace. ■

The motivation of this work is to find necessary and sufficient conditions to guarantee that if a set of incomplete vectors from a union of subspaces fits in an r-dimensional subspace, it is because the set of all their complete yet unknown counterparts indeed lies in an r-dimensional subspace.

2 Introduction

Imagine that an r-dimensional subspace S fits a set of incomplete vectors $\hat{\Xi}$. We want to make sure that the set of all their complete yet unknown counterparts indeed lies in S.

Using the same idea as if the vectors were complete, imagine we had an *extra*, *generic* complete validating vector $\chi \in S^\star \in \mathcal{S}^\star$ that fit in S. It is easy to see that if S fits χ, it is because $S = S^\star$. Furthermore, since S fits $\hat{\Xi}$, this implies that $S^\star$ fits $\hat{\Xi}$.

It is also easy to see that if the subspaces in $\mathcal{S}^\star$ keep no relation with each other, $\hat{\Xi}$ can only fit in one of the subspaces of $\mathcal{S}^\star$ if all its complete counterparts indeed lie in such subspace. This way, S fitting χ would directly imply that all the complete counterparts of $\hat{\Xi}$ indeed lie in S. All the more, it would imply that $S \in \mathcal{S}^\star$.

This is all very nice, but it relies on the fantasy that we had the extra, generic, complete vector χ. Of course, we cannot assume that we have such complete vector. But what if we had several incomplete ones instead? Could a set of extra generic incomplete vectors $\hat{\mathcal{X}}$ *behave* just as χ, allowing us to say that *if* S fits such set, then all the complete counterparts of $\hat{\Xi}$ indeed lie in S?

The answer to this question is yes, and this is precisely what we characterize: when will a set of incomplete vectors $\hat{\mathcal{X}}$ behave as a complete one. This characterization is given in Theorem 1, the main result of the document, which, intuitively, and in a nutshell states that:

> $\hat{\mathcal{X}}$ behaves as a complete vector iff $\hat{\mathcal{X}}$ contains $d-r+1$ vectors such that for every strict subset of n of such vectors, there are at least $n+r$ distinct observed *rows*.

This characterization allows us to fulfill the main task we pursue: determine if a set of incomplete vectors really lies in an r-dimensional subspace whenever it fits in an r-dimensional subspace.

Example 3. Suppose $r = 2$ and

$$\hat{\mathcal{X}} = \begin{bmatrix} 1 & \cdot & \cdot & 3 & 3 \\ 1 & 2 & \cdot & \cdot & 4 \\ 1 & 3 & 4 & \cdot & \cdot \\ \cdot & 4 & 5 & 9 & 6 \\ \cdot & \cdot & 6 & 11 & \cdot \end{bmatrix}.$$

Take the set of the first $d - r + 1 = 4$ vectors. We can verify that every one of its subsets has at least $n + r$ distinct rows with at least one observed entry. For example, if we take the first $n = 2$ vectors, the number of distinct observed rows is 4, which is equal to $n + r$.

We thus conclude that if S fits $\hat{\Xi}$ and $\hat{\mathcal{X}}$, then all the complete counterparts of both $\hat{\mathcal{X}}$ and $\hat{\Xi}$ indeed lie in S. ■

2.1 Insight

There are two fundamental reasons why S fitting a generic χ implies that the complete counterparts of $\hat{\Xi}$ indeed lie in S:

(a) There is only one r-dimensional subspace that fits χ.

(b) Obvious, but essential: χ lies in one and only one of the subspaces of $\mathcal{S}^\star$.

If a set of incomplete vectors $\hat{\mathcal{X}}$ satisfied analogous properties, it would behave just as the complete vector χ in the sense that we would be able to conclude that $S = S^\star \in \mathcal{S}^\star$ if S fits $\hat{\mathcal{X}}$, and the remainder —that $\hat{\Xi}$ indeed lies in S— would follow just as before.

On the other hand, if $\hat{\mathcal{X}}$ failed to have either property, it would fail to behave as the complete vector χ. More precisely, if $\hat{\mathcal{X}}$ fails to satisfy **(a)**, it is evident that we cannot conclude that $S \in \mathcal{S}^\star$; if $\hat{\mathcal{X}}$ fails to satisfy **(b)**, even if there is only one r-dimensional subspace that fits $\hat{\mathcal{X}}$, such subspace might not belong to $\mathcal{S}^\star$, i.e., it could be a *false* subspace. For an example of how this could happen, take Example 2.

In other words, the analogous properties of **(a)** and **(b)** are necessary and sufficient for $\hat{\mathcal{X}}$ to behave as a complete vector. This is precisely what we need to discover: when will a set of generic incomplete vectors $\hat{\mathcal{X}}$ satisfy these two analogous properties.

Remark 1. Observe that **(b)** is substantially different in the complete and incomplete vectors cases. In the complete case, since $\chi \in \mathcal{S}^\star$, and χ is only one vector, we can automatically conclude that almost surely, χ will lie in only one of the subspaces of $\mathcal{S}^\star$.

On the other hand, in the incomplete case, $\hat{\mathcal{X}} \in \mathcal{S}^\star$ does not imply that $\hat{\mathcal{X}}$ lies in one and only one of the subspaces of $\mathcal{S}^\star$, as different vectors from $\hat{\mathcal{X}}$ could belong to different subspaces from $\mathcal{S}^\star$. ■

2.2 The essence

We will see in §4.6 that whether or not S fits a generic $\hat{\mathcal{X}}$ depends only on the position of the observed entries of $\hat{\mathcal{X}}$, namely $\mathbf{\Omega}$. We will also see in §8 that whether a generic $\hat{\mathcal{X}} \in \mathcal{S}^\star$ fits in a single $S^\star \in \mathcal{S}^\star$ or not can also be deduced from $\mathbf{\Omega}$ alone.

Therefore, we may focus on finding conditions on $\mathbf{\Omega}$ to determine when a generic $\hat{\mathcal{X}}$ satisfies the analogous properties **(a)** and **(b)**. This is exactly what we do. Explicitly:

> *We derive deterministic necessary and sufficient conditions on* $\mathbf{\Omega}$ *to* guarantee *that* if *there exists an* r*-dimensional subspace that fits a generic* $\hat{\mathcal{X}}$*, such subspace is unique, and it is because all the vectors of* $\hat{\mathcal{X}}$ *indeed lie in the same subspace of* $\mathcal{S}^\star$*.*

To be clear, the conditions are sufficient in the sense that if $\mathbf{\Omega}$ satisfies such conditions and there exists an r-dimensional subspace that fits a generic $\hat{\mathcal{X}}$, such

subspace is unique, and it must be true that all the vectors of $\hat{\mathcal{X}}$ indeed lie in the same subspace of $\mathcal{S}^\star$. Conversely, the conditions are necessary in the sense that even if there exists an r-dimensional subspace that fits a generic $\hat{\mathcal{X}}$, if $\mathbf{\Omega}$ does not satisfy such conditions, such subspace may not be unique, and it cannot be guaranteed that the vectors of $\hat{\mathcal{X}}$ lie in the same subspace of $\mathcal{S}^\star$.

The conditions to guarantee that all the elements of $\hat{\mathcal{X}}$ indeed lie in the same subspace of $\mathcal{S}^\star$ are given in Theorem 1. They imply and rely on the conditions for uniqueness, which are given in Theorem 2. As we could see in §2, these conditions are extremely simple and concrete, and depend only on the most elemental invariants of $\mathbf{\Omega}$: essentially, cardinalities of its subsets. Both of these results, the main ones of the document, are presented formally in §6, the section of results. Together, they characterize when a set of incomplete vectors behaves as a complete one, which allows us to verify our final goal: when $\hat{\Xi}$ indeed lies in S.

2.3 Organization of the document

In §3 it is given a brief discussion about previous and related work; this helps as preamble to give some motivation for this problem and talk about some particularly interesting applications of this work that give simple yet powerful consequences of its results. In §4 it is given a detailed exposition of the setup that will be used in the remainder of the document. In §5 the assumptions of this work are stated, explained and discussed. The main results are given in §6.

The analysis to prove Theorem 2 is presented in §7, and the one to prove Theorem 1 in §8. In §9 it is offered an intuitive explanation of the key ideas of the results and the assumptions are discussed in more detail, as well as some simple generalizations. Finally, a brief proposal for future research is given in §10.

To make the reading of this document easier, the main symbols, terms, statements, definitions, examples, etc., are referenced in the whole document in its electronic version; alternatively, an index and al list of symbols is also provided at the end of the document.

3 Preamble

With the arrival of big data come big challenges: we want to find useful information in our datasets quickly, cleverly, using as few resources as possible. Fortunately, in uncountable applications we may use subspaces to model our data, and this greatly simplifies things.

But this is not it. As if finding useful information quickly, cleverly and efficiently were not ambitious enough endeavors, we also want —and many times, need— to achieve these tasks only with partial information, which comes as no surprise, as the bigger our data, the more likely it is incomplete.

Fortunately, subspaces have a natural way of handling missing data, as data in subspaces have certain structure, and that gives us a way to infer the missing entries. The problem of handling missing data has attracted a lot of attention in recent years. Remarkable work has been done to identify a subspace that fits certain incomplete data, e.g., [2], to detect if an incomplete datum fits in a certain subspace[3], or even to do subspace clustering from missing data[4], but the converse problem, in the sense we discuss in **M1**, has been left unattended, and remained, to the best of our knowledge, an open problem until now.

But again, as if the task of finding useful information quickly, cleverly, efficiently, and only from partial information were not bold enough, we also want something else. We want to make sure that if we reach a conclusion from our data, such conclusion is correct. In other words, we want to make sure that the information we found is not a product of chance; the larger our data, the more outliers, the more likely we will find *something*, but that doesn't mean that that *something* is *true*. If we toss a coin a trillion times, we will very likely see many sequences of many heads in a row, but that doesn't mean that a sequence of many heads in a row is very likely. The more data we have, the more likely *some* subspace will *fit some* of it, but that doesn't mean that our data really *lies* in such subspace.

These is precisely the task that we are interested on: how to determine when certain incomplete data really *lies* in a subspace whenever it *fits* in a subspace.

Notice the subtle but fundamental difference between this work and, for example, the matched subspace detection with missing data problem in [3], where they are concerned with determining if an incomplete datum *fits* a subspace, using only information about such datum and the subspace. Here we are given an incomplete dataset that we already know *fits* in a subspace, and we want to make sure that it really *lies* in it, using the dataset as a whole, exploiting information about the relation between their datums. Similarly, in [2] they are concerned with identifying a subspace that *fits* certain incomplete data, under the assumption that the data *lies* in a subspace. Here we drop such assumption; we are given an incomplete dataset and a subspace that *fits* it, and we want to know if the dataset really *lies* in such subspace.

The problem of determining if certain incomplete data really *lies* in a subspace whenever it *fits* in a subspace is tightly related to the problem of identifying when there is only one subspace that fits such data. We answer these questions by characterizing when and only when a set of incomplete vectors

behaves as a single but complete one, in the sense described in §2.

Being these so fundamental problems, answering these questions should be enough motivation by itself, as they essentially apply to virtually every problem involving subspaces and missing data. Nevertheless, just for completeness, we mention just a few motivating applications, to give an idea of the scope and relevance of these results.

M1. Consider the low-rank matrix completion problem[2]: given that all the columns of a matrix Ξ lie in the same r-dimensional subspace S, under what conditions is S the only r-dimensional subspace that fits a subset of the entries of such matrix, $\hat{\Xi}$?

The condition that Ξ lies in the same r-dimensional subspace trivially implies that there exists an r-dimensional subspace that fits $\hat{\Xi}$. As we explained in §1, the converse is not necessarily true (see Examples 1, 17 and §8 for a more detailed explanation).

Our work immediately provides a **converse** of the low-rank matrix completion problem: all the columns of a matrix Ξ lie in the same r-dimensional subspace if there exists an r-dimensional subspace that fits $\hat{\Xi}$ and an additional generic $\hat{\mathfrak{X}}$ observed in a set $\mathbf{\Omega}$ satisfying the conditions of Theorem 1.

M2. Under the same setup of low-rank matrix completion, most algorithms, e.g., nuclear norm minimization[2], detect an r-dimensional subspace S that fits an incomplete dataset $\hat{\Xi}$, and claim that with high probability, the detected subspace S is the only r-dimensional one that does. Theorem 2 provides a deterministic validation check for any such algorithm: S is almost surely the unique r-dimensional subspace that fits $\hat{\Xi}$ if in addition it also fits a generic $\hat{\mathfrak{X}}$ observed in a set $\mathbf{\Omega}$ satisfying the conditions of Theorem 2.

As we said in **M1**, under this setup Ξ is already assumed to belong to the same subspace, so Theorem 1 is not even required here; Theorem 2 alone is sufficient for the purposes of this problem.

M3. Extending **M2**, there is no reason to stop with low-rank matrix completion. Theorem 1 provides a deterministic validation check for *any* algorithm that performs low-rank, or even high-rank matrix completion[5], or any algorithm that finds a subspace that fits data, e.g., the EM algorithm derived in [4].

M4. Continuing with **M3**, a universal deterministic validation check on the output of any algorithm opens the door to answering an important open question: the real sample complexity of subspace clustering with missing data[4], which is somewhat equivalent to the sample complexity of high-rank matrix completion[5]. One can see in [4] that the gist of this problem is to be able to identify *false* subspaces that for some unfortunate circumstances *could* deceivingly appear to fit certain data.

M5 Of course, sometimes even when we know that our data lies in a subspace — or want to approximate it with a subspace anyway— we don't always know the dimension of such subspace —or the minimum possible dimension of a subspace that approximates it nicely. Theorem 1 can be used iteratively to find with certainty the lowest-dimensional subspace or the *minimal* union of subspaces that fit certain data.

Not pretending to do a survey on the applications of subspaces, we think these motivations should be enough to give an idea of the scope and power of these results. With this, we move on.

4 Setup

In this section we fully describe the setup and notation that will be used in the remainder of the document.

4.1 Subspaces and bases

Let $\mathcal{S}^\star = \{S_k^\star\}_{k=1}^K$ be a set of K distinct r-dimensional subspaces of $\mathbb{R}^d$.

We use $S^\star$ to denote an arbitrary subspace from $\mathcal{S}^\star$, and $U^\star$ to denote a basis of $S^\star$, i.e., whenever possible, we drop the subscript k, which is generally used to index subspaces, and unless otherwise stated, runs from 1 to K.

Example 4. Let $d = 5$, $r = 2$, and

$$U^\star = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{bmatrix}.$$

Then $S^\star = \text{span}\{U^\star\}$. ■

In general, we use S to denote an arbitrary subspace and U to denote one of its bases.

4.2 Vectors and bases

Let $\mathcal{X} := \{x_i\}_{i=1}^N$ be denote a collection of vectors of $\mathbb{R}^d$ that lies in the union of the subspaces of $\mathcal{S}^\star$.

As we said before, each x_i is assumed to lie in one of the subspaces of $\mathcal{S}^\star$. This correspondence is described by $\mathcal{K} = \{k_i\}_{i=1}^N$, a multiset of indices in $\{1, ..., K\}$ that specifies that x_i lies in $S_{k_i}^\star$. To keep notation from getting out of hand, we use $S_i^\star$ as shorthands for $S_{k_i}^\star$, and $U_i^\star$ to denote a basis of $S_i^\star$.

We use x to denote an arbitrary element of $\mathcal{X}$ that lies in $S^\star$, i.e., whenever possible, we drop the subscript i, which is generally used to index vectors, and unless otherwise stated, runs from 1 to N.

Example 5. With the same setup as in Example 4, let $N = 3$ and

$$x_1 = \begin{bmatrix} 2 \\ 2 \\ 2 \\ 2 \\ 2 \end{bmatrix}, \quad x_2 = \begin{bmatrix} 3 \\ 6 \\ 9 \\ 12 \\ 15 \end{bmatrix}, \quad x_3 = \begin{bmatrix} 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{bmatrix}.$$

Then

$$\mathcal{X} = \{x_1, x_2, x_3\} = \begin{bmatrix} 2 & 3 & 2 \\ 2 & 6 & 3 \\ 2 & 9 & 4 \\ 2 & 12 & 5 \\ 2 & 15 & 6 \end{bmatrix}.$$

It is easy to see that x_1, x_2 and x_3 belong to $S^\star$, i.e., $k_1 = k_2 = k_3$. ■

4.3 Observation sets

We are interested on partially observed vectors. We handle this with $\mathbf{\Omega} := \{\omega_i\}_{i=1}^N$, a set of N sets that specifies that x_i is only observed in the positions of the set ω_i. Since $x_i \in \mathbb{R}^d$, $\omega_i \subset \{1, ..., d\}$.

We make two assumptions about the entries we observe, only to simplify the analysis, but these can be most easily generalized:

(i) $|\omega_i| = r + 1$ for every i.

(ii) $\bigcup_i \omega_i = \{1, ..., d\}$.

We use ω to denote an arbitrary subset of $\{1, ..., d\}$ of size $r + 1$, i.e., whenever possible we drop the subscript i.

We also use Ω, $\breve{\Omega}$ and $\bar{\Omega}$ to denote arbitrary collections of sets of ω's; typically subsets of $\mathbf{\Omega}$.

Definition 1 (n, m)**.** *Given Ω (resp. $\breve{\Omega}$ and $\bar{\Omega}$), we define n and m (resp. $\breve{n}$, $\breve{m}$ and $\bar{n}$, $\bar{m}$) as:*

$$n := |\Omega|,$$
$$m := |\bigcup_{\omega \in \Omega} \omega|,$$

Definition 2. *Given $\Omega \subset \mathbf{\Omega}$, we define $\mathcal{I} := \{i : \omega_i \in \Omega\}$ and $\mathcal{J} := \bigcup_{\omega \in \Omega} \omega$.*

Observe that $|\Omega|$ is the number of sets that Ω contains, i.e., $|\mathcal{I}|$, for example, $|\mathbf{\Omega}| = N$.

Unless otherwise stated, we use j to index the elements of $\{1, ..., d\}$, and typically to denote that such element belongs to some set ω or to index an entry of a vector, for example, u_j denotes the j^{th} entry of u. This way, intuitively, $\mathcal{J}$ is the set of j's contained in the sets of Ω, m is the number of distinct j's that are contained in the sets of Ω, n is the number of ω's that Ω has, and for $\Omega \subset \mathbf{\Omega}$, $\mathcal{I}$ is the set of i's such that $\omega_i \in \mathbf{\Omega}$ also belongs to Ω.

For convenience, rather than listing the set of sets to specify Ω, we typically use a $d \times n$ matrix whose $(j, i)^{th}$ entry is *observed*, denoted by $\times$, if $j \in \omega_i$, and *missing* otherwise, denoted by $\cdot$. When there is no room for confusion, we use Ω to denote such matrix. Under this convention, $\mathcal{J}$ can be thought of as the set of rows with at least one observed entry, and m as the number of such rows.

Example 6. With the same setup as in Example 5. Let $\omega_1 = \{1, 2, 3\}$, $\omega_2 = \{2, 3, 4\}$ and $\omega_3 = \{3, 4, 5\}$. Then

$$\mathbf{\Omega} = \{\omega_1, \omega_2, \omega_3\} = \begin{bmatrix} \times & \cdot & \cdot \\ \times & \times & \cdot \\ \times & \times & \times \\ \cdot & \times & \times \\ \cdot & \cdot & \times \end{bmatrix}.$$

If we let $\Omega = \{\omega_1, \omega_2\}$, then $\mathcal{I} = \{1, 2\}$, $\mathcal{J} = \{1, 2, 3, 4\}$, $n = 2$ and $m = 4$. ■

4.4 Incomplete vectors, bases and subspaces

We are now ready to define incomplete vectors.

Definition 3 ($\hat{x}$)**.** *Given ω, we define $\hat{x}$ as the vector with d components whose j^{th} entry is equal to the j^{th} entry of x if $j \in \omega$, and otherwise has a value of* missing*, denoted by $\cdot$ (resp. for ω_i and $\hat{x}_i$).*

Notice that $\hat{x}$ depends on ω. Technically, we could specify this by writing $\hat{x}_\omega$, but the index ω is redundant, and we want our notation to be as simple as possible. For a collection of vectors we simply have $\hat{\mathcal{X}} := \{\hat{x}_i\}_{i=1}^N$. When there is no room for confusion, we equivalently use $\mathcal{X}$ and $\hat{\mathcal{X}}$ to denote the $d \times N$ matrices with $\{x_i\}_{i=1}^N$ and $\{\hat{x}_i\}_{i=1}^N$ as its columns.

Definition 4 (x_ω)**.** *We define x_ω as the vector in $\mathbb{R}^{|\omega|}$ whose entries are equal to the observed entries of $\hat{x}$.*

For subspaces we have something similar.

Definition 5 ($\mathbb{R}^d_\omega$, $\hat{S}$ and $\hat{U}$)**.** *Let $\mathbb{R}^d_\omega$ be the span of the canonical vectors of $\mathbb{R}^d$ corresponding to the elements of ω. We define $\hat{S}$ as the projection of S onto $\mathbb{R}^d_\omega$, and $\hat{U}$ as the $d \times r$ matrix with the entries of U in the positions of ω, and zeros elsewhere (resp. for ω_i, $\hat{S}_i$ and $\hat{U}_i$).*

It is easy to see that $\text{span}\{\hat{U}\} = \hat{S}$. Conversely, the rows of any basis of $\hat{S}$ must be zero in the positions that don't belong to ω.

Similar to $\hat{x}$, $\hat{S}$ and $\hat{U}$ depend on ω. Technically, we could specify this by writing $\hat{S}_\omega$ or $\hat{U}_\omega$, but the index ω is redundant, and we want to keep our notation from getting out of hand. For this same purpose, we use $\hat{S}^\star_i$ as shorthand for $\hat{S}^\star_{k_{i_{\omega_i}}}$, and similarly for $\hat{U}^\star_i$ and $\hat{U}^\star_{k_{i_{\omega_i}}}$.

Definition 6 (U_ω and S_ω)**.** *Given ω, we define U_ω as the $|\omega| \times r$ matrix with the non-zero rows of rows of $\hat{U}$, and $S_\omega := \text{span}\{U_\omega\}$.*

To simplify our notation, we use $S^\star_{\omega_i}$ as shorthand for $S^\star_{k_{i_{\omega_i}}}$, and similarly for $U^\star_{\omega_i}$.

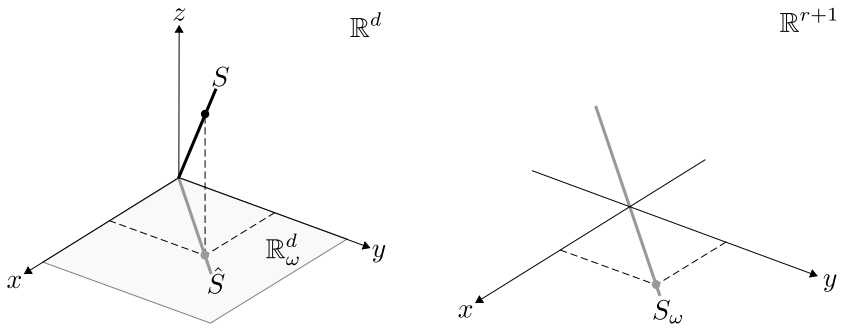

(a) $\hat{S}$ is the result of projecting S onto $\mathbb{R}^d_\omega$.

(b) S_ω is the *restriction* of S to the positions of ω; a subspace in $\mathbb{R}^{r+1}$.

Figure 2: In this example, $r = 1$, and $\omega = \{1, 2\}$, so $\mathbb{R}^d_\omega$ is the (x, y)-plane. $\hat{S}$ is a line in $\mathbb{R}^3$ that lies in the (x, y)-plane, while S_ω is a line in $\mathbb{R}^2$.

Example 7. With the same setup as in Examples 4, 5 and 6, we would obtain

$$\hat{x}_1 = \begin{bmatrix} 2 \\ 2 \\ 2 \\ \cdot \\ \cdot \end{bmatrix}, \ \hat{x}_2 = \begin{bmatrix} \cdot \\ 6 \\ 9 \\ 12 \\ \cdot \end{bmatrix}, \ \hat{x}_3 = \begin{bmatrix} \cdot \\ \cdot \\ 4 \\ 5 \\ 6 \end{bmatrix}, \quad \hat{\mathcal{X}} = \begin{bmatrix} 2 & \cdot & \cdot \\ 2 & 6 & \cdot \\ 2 & 9 & 4 \\ \cdot & 12 & 5 \\ \cdot & \cdot & 6 \end{bmatrix},$$

$$x_{\omega_1} = \begin{bmatrix} 2 \\ 2 \\ 2 \end{bmatrix}, \ x_{\omega_2} = \begin{bmatrix} 6 \\ 9 \\ 12 \end{bmatrix}, \ x_{\omega_3} = \begin{bmatrix} 4 \\ 5 \\ 6 \end{bmatrix},$$

$$\hat{U}_1 = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \\ 0 & 0 \\ 0 & 0 \end{bmatrix}, \quad \hat{U}_2 = \begin{bmatrix} 0 & 0 \\ 1 & 2 \\ 1 & 3 \\ 1 & 4 \\ 0 & 0 \end{bmatrix}, \quad \hat{U}_3 = \begin{bmatrix} 0 & 0 \\ 0 & 0 \\ 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{bmatrix},$$

$$U_{\omega_1} = \begin{bmatrix} 1 & 1 \\ 1 & 2 \\ 1 & 3 \end{bmatrix}, \quad U_{\omega_2} = \begin{bmatrix} 1 & 2 \\ 1 & 3 \\ 1 & 4 \end{bmatrix}, \quad U_{\omega_3} = \begin{bmatrix} 1 & 3 \\ 1 & 4 \\ 1 & 5 \end{bmatrix}.$$

■

Before we move to other things, one technical definition that will simplify our argumentation greatly without loss of generality. We will discuss more about this in §9.3.

Definition 7 (Degenerate subspace)**.** *We say an r-dimensional subspace is* degenerate *iff there exists an $\upsilon \subset \{1, ..., d\}$ with $|\upsilon| \leq r$, such that $\dim S_\mathcal{U} < |\upsilon|$.*

Remark 2. Definition 7 is saying that a subspace is non-degenerate iff for every $\upsilon \subset \{1, ..., d\}$ with $|\upsilon| \leq r$, $\dim S_\mathcal{U} = |\upsilon|$, i.e. $S_\mathcal{U} = \mathbb{R}^{|\mathcal{U}|}$, or equivalently, iff

every $|\upsilon| \times r$ matrix formed with the rows of any of its bases is full-rank. Notice that almost every subspace is non-degenerate. ■

Example 8. Unless otherwise stated, subspaces of all examples in the document are non-degenerate. Here is an example of a degenerate one:

$$U = \begin{bmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 1 \end{bmatrix}.$$

If we take $\upsilon = \{2, 3\}$, one can verify that $\dim S_\upsilon = 1 < 2 = |\upsilon|$ by simply looking at the bottom 2×2 minor of U, which is rank-defficient. ■

4.5 Fitting incomplete vectors

Let us now define what it means to *fit* an incomplete vector.

Definition 8 (To fit $\hat{x}$)**.** *We say S* fits *$\hat{x}$ iff there exists a vector in S that is equal to $\hat{x}$ in all its observed entries, or equivalently, iff $x_\omega \in S_\omega$.*

Definition 9 (To fit $\hat{\mathcal{X}}$)**.** *We say S* fits *$\hat{\mathcal{X}}$ iff S fits $\hat{x}_i$ for every i.*

Notice that each $\hat{x}_i$ might belong to a different subspace in $\mathcal{S}^\star$.

Example 9. With the same setup as Examples 1 and 2, it is easy to see that S fits $\hat{\mathcal{X}}$. ■

4.6 Fitting generic vectors ∼ fitting observation sets

We now formalize what we mean by fitting generic vectors. Intuitively, when we say that a vector is generic, we mean that it could be *any* vector, and whenever we say that S fits a generic vector from $S^\star$, what we formally mean is that S fits *every* vector from $S^\star$.

The same for an incomplete vector. Moreover, with Definition 8, it is easy to see that whether or not S fits *every* $\hat{x} \in S^\star$ depends only on ω. Therefore, whenever we informally say that S fits a generic $\hat{x}$, what we formally mean is that S fits ω in the following sense.

Definition 10 (To fit ω)**.** *Given $S^\star$, we say that S* fits *ω iff S_ω fits* every *$x_\omega \in S^\star_\omega$.*

Remark 3. Notice that S will fit ω iff $S^\star_\omega \subset S_\omega$, i.e., iff for every $x \in S^\star$ there exists a $u \in S$ such that $u_\omega = x_\omega$.

In general, S *fits* ω iff $S^\star_\omega \subset S_\omega$. Nevertheless since $S^\star$ is non-degenerate, $S^\star_\omega$ is an r-dimensional subspace, so whenever S is also an r-dimensional subspace, we also have $S_\omega \subset S^\star_\omega$, whence S *fits* ω iff $S_\omega = S^\star_\omega$. ■

Example 10. Let $r = 1$ and suppose

$$U = \begin{bmatrix} 2 \\ 2 \\ 3 \end{bmatrix}, \quad U^\star = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \quad \omega = \begin{bmatrix} \times \\ \times \\ \cdot \end{bmatrix}.$$

Since S would fit any $\hat{x} \in S^\star$, we say S fits ω. ■

Similarly, when we say that a set of vectors $\mathcal{X} \in \mathcal{S}^\star$, consisting of $x_1 \in S_1^\star, ..., x_N \in S_N^\star$, is generic, we mean that x_1 could be *any* vector from $S_1^\star$, x_2 could be *any* vector from $S_2^\star$, and so on. Whenever we say that S fits a generic set of vectors $\mathcal{X} \in \mathcal{S}^\star$, what we formally mean is that S fits *every* $x_i \in S_i^\star$ for every i.

The same for sets of incomplete vectors. With Definition 9 it is easy to see that whether or not S fits *every* $\hat{\mathcal{X}} \in \mathcal{S}^\star$ depends only on $\mathbf{\Omega}$. Therefore, whenever we informally say that S fits a generic $\hat{\mathcal{X}}$, what we formally mean is that S fits $\mathbf{\Omega}$, with the following.

Definition 11 (To fit Ω)**.** *Given $\mathcal{S}^\star$ and $\mathcal{K}$, we say that S* fits *Ω iff S fits $\hat{x}_i$ for* every *$x_i \in S_i^\star$ and* every *$i \in \mathcal{I}$.*

Remark 4. Recall that $S_i^\star$ is a shorthand for $S_{k_i}^\star$ and k_i is the index in $\mathcal{K}$ that specifies that x_i lies in $S_{k_i}^\star$. Hence the dependency on $\mathcal{K}$ in Definition 11. ■

Notice that there are many equivalent ways of defining what it means to fit Ω. For example, we could also say that S *fits* Ω iff S_{ω_i} fits x_{ω_i} for *every* $x_{\omega_i} \in S_{\omega_i}^\star$ and *every* $i \in \mathcal{I}$, or we could define it as in the next Remark.

Remark 5. In general, S *fits* $\mathbf{\Omega}$ iff $S_{\omega_i}^\star \subset S_{\omega_i}$ for every i. Nevertheless, since $S^\star$ is non-degenerate, $S_{\omega_i}^\star$ is an r-dimensional subspace, so whenever S is also an r-dimensional subspace, we also have that $S_{\omega_i} \subset S_{\omega_i}^\star$ for every i, whence S *fits* $\mathbf{\Omega}$ iff $S_{\omega_i} = S_{\omega_i}^\star$ for every i. ■

Example 11. Suppose $r = 1$ and

$$U = \begin{bmatrix} 2 \\ 2 \\ 3 \end{bmatrix}, \quad U_1^\star = \begin{bmatrix} 1 \\ 1 \\ 1 \end{bmatrix}, \quad U_2^\star = \begin{bmatrix} 1 \\ 2 \\ 3 \end{bmatrix}, \quad \mathbf{\Omega} = \begin{bmatrix} \times & \times \\ \times & \cdot \\ \cdot & \times \end{bmatrix}.$$

Since S would fit every $\hat{x}_1 \in S_1^\star$ *and* every $\hat{x}_2 \in S_2^\star$, we say S fits $\mathbf{\Omega}$. ■

Remark 6. When we informally say that an r-dimensional subspace S fits a *generic* $\hat{x}$, what we formally mean is that S fits ω, i.e., that S fits $\hat{x}$ for *every* $x \in S^\star$. This guarantees that S is somehow independent of one particular instance of $\hat{x}$. This is essential for our analysis, because if S is to fit ω, then S_ω must fit x_ω for *every* $x_\omega \in S_\omega^\star$. This implies that S must satisfy $S_\omega = S_\omega^\star$ (see Remark 3).

In contrast, S need not satisfy this to fit *one particular* $\hat{x}$. For instance, S could fit *one particular* $\hat{x}$ by just fixing the observed entries of $\hat{x}$ in the positions

of ω of a spanning vector of S, e.g., with the same setup as in Example 4, let

$$x = \begin{bmatrix} 2 \\ 3 \\ 4 \\ 5 \\ 6 \end{bmatrix}, \qquad \omega = \begin{bmatrix} \times \\ \times \\ \times \\ \cdot \\ \cdot \end{bmatrix}, \qquad \hat{x} = \begin{bmatrix} 2 \\ 3 \\ 4 \\ \cdot \\ \cdot \end{bmatrix}.$$

Then we could construct

$$U = \begin{bmatrix} 2 & u_{12} \\ 3 & u_{22} \\ 4 & u_{32} \\ u_{41} & u_{42} \\ u_{51} & u_{52} \end{bmatrix},$$

and U would fit $\hat{x}$ for any choices of u_{12}, u_{22} and u_{32}, so S would *need not* satisfy $S_\omega = S^\star_\omega$ to fit $\hat{x}$.

Similarly, when we informally say that an r-dimensional subspace S fits a *generic* $\hat{\mathfrak{X}}$, what we formally mean is that S fits $\mathbf{\Omega}$, i.e., that S fits $\hat{x}_i$ for *every* $x_i \in S^\star_i$ *and every* i. This guarantees that S is somehow independent of one particular instance of $\hat{\mathfrak{X}}$. This is essential for our analysis, because S must satisfy $S_{\omega_i} = S^\star_{\omega_i}$ for every i in order to fit $\mathbf{\Omega}$.

■

5 Assumptions

Now that we have fully specified our setup, we use this section to give a detailed, precise and unified list of our assumptions, in order to present them all together and clearly, to emphasize how lenient they are, how they are mostly used to ease our argumentation, and how easily they can be generalized.

A1. $x_i \in S_i^\star$ for every i. This is just the basic setup of the problem: that our data lie in a union of subspaces. This is to give some structure to $\mathcal{X}$. Without this assumption, we cannot possibly hope to determine the missing values in $\hat{\mathcal{X}}$, as they could be arbitrary, whence nothing can be said about the complete counterpart $\mathcal{X}$, and no low-dimensional subspace can be guaranteed to fit $\mathcal{X}$.

A2. All subspaces in $\mathcal{S}^\star$ are assumed to be r-dimensional subspaces of $\mathbb{R}^d$. This is just to simplify our arguments, but can be easily generalized to the case where the subspaces in $\mathcal{S}^\star$ are of different dimensions (see §9.3). Observe that we are not assuming anything about K. For all we know, K could be arbitrarily large; even larger than N.

A3. All subspaces in $\mathcal{S}^\star$ are assumed to be non-degenerate. This is just to simplify our arguments, but can be easily generalized, *if necessary*. We emphasize, *if necessary*, because fortunately, the set of degenerate subspaces has measure zero, thus our results hold, without any further modification, for almost every $\mathcal{S}^\star$. For a further discussion about degenerate subspaces see §9.3.

A4. $|\omega_i| = r + 1$ for every i. Observe that if $|\omega| \leq r$, there is no possible way to determine the subspace where x really lies. More precisely, since subspaces in $\mathcal{S}^\star$ are non-degenerate, a vector observed in fewer than $r+1$ entries could belong to any of the subspaces in $\mathcal{S}^\star$.

On the other hand, if $|\omega| > r + 1$ it can only be easier to determine if x really belongs to S, as any subspace that fits it will have to satisfy more restrictions. In other words, it is harder to fit $\hat{x}_1$ than to fit $\hat{x}_2$ if $|\omega_1| > |\omega_2|$. This assumption can be immediately generalized using this simple observation.

This way, rather than an assumption —being $|\omega| > r$ a requirement for the task that we want to achieve— this is just a convenience statement to simplify our arguments, analysis and notation, that at the same time states that we are working under the most minimal assumptions on $|\omega|$. More about this is discussed in §9.3.

Notice that we assume nothing about $\mathbf{\Omega}$ being spread uniformly, at random, or anything of the sort; our results apply to completely arbitrary $\mathbf{\Omega}$'s.

A5. $\mathcal{J} := \bigcup_i \omega_i = \{1, ..., d\}$. This is just for simplicity of notation and argumentation. Observe that if $\mathbf{\Omega}$ has a fully unobserved row, there will

always be infinitely many subspaces that fit $\mathbf{\Omega}$. We can only determine when the projection of such subspaces onto $\mathbb{R}^d_{\boldsymbol{\mathcal{J}}}$ will be unique, i.e., when the restriction of a subspace that fits $\mathbf{\Omega}$, is unique. We can easily generalize this by working only with the observed entries of $\mathbf{\Omega}$, and alternatively defining $\{1, ..., d\} := \boldsymbol{\mathcal{J}}$, as we can say nothing anyway about the entries where no row is observed.

A6. There exists a non-degenerate r-dimensional subspace that fits $\mathbf{\Omega}$. Rather than an assumption, this is the motivation of the document: assuming that there is a subspace that fits certain data, we want to determine when such data really *lies* in a subspace. Since we are assuming that subspaces in $\mathcal{S}^\star$ are non-degenerate, we know that if our data really lies in a subspace, it is a non-degenerate one. If certain degenerate subspace fits our data, we would already know our data does not really lies in it. We use this assumption mainly to avoid all this uninteresting argumentation in every statement.

6 Results

To avoid being anal in every statement of the document without giving up being precise, we would like to give an important remark about our results before presenting them:

All our results hold under assumptions A1-A6 above, and for almost every $\mathcal{S}^\star$ under the natural Lebesgue measure.

In other words, our results hold for almost every dataset lying in a union of subspaces.

Having said that, we are now ready to present our results. We would like to take advantage of this section to also give some intuitive meaning to them, to emphasize how simple they are, and to give an example of their usage.

Recall that the main goal is to determine how can we make sure that if certain data *fit* in a subspace S, it is because such data really *lie* in a subspace.

As discussed in §2, we can answer this question by characterizing when a set of incomplete vectors observed only in $\mathbf{\Omega}$, behaves as a complete one. Such characterization, intuitively stated in §2 is formalized by our main result:

Theorem 1 (All of a kind)**.** *Suppose an r-dimensional subspace S fits $\mathbf{\Omega}$, and that there exists an $\breve{\Omega} \subset \mathbf{\Omega}$ of size $d - r + 1$ such that*

$$m \geq n + r \qquad \forall\ \Omega \subsetneq \breve{\Omega}. \tag{1}$$

Then $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$. Furthermore, S is the only r-dimensional subspace that fits $\mathbf{\Omega}$, and $S \in \mathcal{S}^\star$.

Conversely, if no such $\breve{\Omega}$ exists, k_i might *be different from $k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$, there* could *be infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$, and even if there is only one such subspace, it* might *not even belong to $\mathcal{S}^\star$.*

Before moving on, we give some intuitive interpretations of Theorem 1:

(a) That there is only one subspace that fits $\mathbf{\Omega}$ means that there is only one subspace that fits a set of generic incomplete vectors $\hat{\mathfrak{X}}$.

(b) That $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$ means that all the vectors of $\hat{\mathfrak{X}}$ lie in the same subspace of $\mathcal{S}^\star$.

In other words, Theorem 1 is telling us precisely what we wanted: when will a generic $\hat{\mathfrak{X}}$ satisfy the desired properties **(a)** and **(b)** from §2.1, and hence when it will behave as a complete vector.

But not only that. Notice that the converse is telling us that if $\mathbf{\Omega}$ does not satisfy the conditions of the theorem, then it cannot be guaranteed that all the

vectors of $\hat{\mathcal{X}}$ lie in the same subspace of $\mathcal{S}^\star$, whence $\hat{\mathcal{X}}$ will not behave as a complete vector.

In conclusion, Theorem 1 is telling us that a generic $\hat{\mathcal{X}}$ will behave as a complete generic vector iff $\mathbf{\Omega}$ satisfies the condition of the theorem, namely that it contains a set $\breve{\Omega}$ of size $d - r + 1$ that satisfies (1). One simple and intuitive interpretation of (1) is that

> *For every strict subset of* $\breve{\Omega}$ *with* n columns*, the number of distinct* rows *with at least one observation,* m*, is at least* $n + r$*.*

Example 12. Continuing with Example 3, we have that

$$\mathbf{\Omega} = \begin{bmatrix} \times & \cdot & \cdot & \times & \times \\ \times & \times & \cdot & \cdot & \times \\ \times & \times & \times & \cdot & \cdot \\ \cdot & \times & \times & \times & \times \\ \cdot & \cdot & \times & \times & \cdot \end{bmatrix}.$$

It is easy to see that $\mathbf{\Omega}$ satisfies the conditions of Theorem 1. Explicitly, take $\breve{\Omega} = \{\omega_1, ..., \omega_4\}$. One may verify that $\breve{\Omega}$ satisfies (1). If S fits $\mathbf{\Omega}$, then $k_i = k_{\breve{\imath}}$ for every $(i, \breve{\imath})$ and $S = S_i^\star$ is the only r-dimensional subspace that fits $\mathbf{\Omega}$.

For an example of the converse statement of Theorem 1, consider the same setup as in Example 17. Observe that there exists no $\breve{\Omega} \subset \mathbf{\Omega}$ that satisfies the conditions of Theorem 1. Thus the columns of $\mathcal{X}$ *might* not belong to the same subspace, which is precisely the case. We can also see that S is none of the subspaces in $\mathcal{S}^\star$. ■

As we mentioned in §2.2, determining when there is only one r-dimensional subspace that fits $\mathbf{\Omega}$ is essential for the proof of Theorem 1. The answer to this is given by our second main result.

Theorem 2 (Characterization of uniqueness)**.** *There is* only one r*-dimensional subspace that fits* $\mathbf{\Omega}$ *iff there exists an* $\breve{\Omega} \subset \mathbf{\Omega}$ *of size* $d - r$ *such that*

$$m \geq n + r \qquad \forall\ \Omega \subset \breve{\Omega}. \tag{2}$$

Notice that the requirement of Theorem 1 is slightly stronger than the requirement of Theorem 2: Theorem 1 requires $\mathbf{\Omega}$ to contain a set $\breve{\Omega}$ of size $d-r+1$ that satisfies (2) for every one of its $d-r$ subsets of size $d-r$, while Theorem 2 only requires one $\breve{\Omega}$ of size $d - r$ that satisfies (2).

In other words, once we know that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$, we only need a little bit more to make sure that all the vectors of $\mathcal{X}$ indeed lie in the same subspace.

Example 13. With the same setup as in Example 11, it is clear that $\mathbf{\Omega}$ is the only possible set of size $d - r$, so take $\breve{\Omega} = \mathbf{\Omega}$. One can trivially verify that $m \geq n + r$ for every $\Omega \subset \mathbf{\Omega}$, so $\mathbf{\Omega}$ satisfies the conditions of Theorem 1, hence there is only one r-dimensional subspace that fits $\mathbf{\Omega}$.

Nevertheless, observe that $\mathbf{\Omega}$ does not satisfy the conditions of Theorem 1, as it does not even have $d - r + 1$ sets.

This is one case where we can guarantee that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$, but we cannot yet guarantee that $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$, nor that the subspace that fits $\mathbf{\Omega}$ belongs to $\mathcal{S}^\star$.

Notice that the difference between this $\mathbf{\Omega}$ and the $\mathbf{\Omega}$ from Example 18, where we can guarantee that $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$, is just one ω. ■

7 Uniqueness

Determining when there is *only one* r-dimensional subspace that fits $\mathbf{\Omega}$ is essential for the proof of Theorem 1. The answer is given by Theorem 2. We will thus start with the proof of the later, and will leave the proof of the former for §8.

7.1 The subspace $\mathcal{U}_\Omega$

With no further ado, we begin our analysis. Consider the set:

$$\mathcal{U}_\Omega := \left\{ u \in \mathbb{R}^d : \bigcap_{i \in \mathcal{I}} \{ u_{\omega_i} \in S^\star_{\omega_i} \} \right\}.$$

Observe that $\mathcal{U}_\Omega$ is a function of $\mathcal{S}^\star$, $\mathcal{K}$ and Ω (see Remark 5).

In order to show Theorem 2 we will prove that $\mathcal{U}_\Omega$ is a subspace, that it contains all the r-dimensional subspaces that fit Ω, that it fits Ω, and will determine its dimension. We will then conclude that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$ iff $\dim \mathcal{U}_{\mathbf{\Omega}} = r$.

We begin our work towards these goals. Similar to the definition of $\mathcal{U}_\Omega$, let

$$\mathcal{U}_\omega := \{ u \in \mathbb{R}^d : u_\omega \in S^\star_\omega \}.$$

In other words, $\mathcal{U}_\omega$ is the set of all u's such that u_ω fits in $S^\star_\omega$. Notice that $\mathcal{U}_\omega$ is a function of $S^\star$ and ω.

Lemma 1. *$\mathcal{U}_\omega$ and $\mathcal{U}_\Omega$ are subspaces.*

Proof. Let $u, v \in \mathcal{U}_\omega$, and $w = \varsigma u + \zeta v$ for some scalars ς, ζ. Since $S^\star_\omega$ is a subspace, $w_\omega \in S^\star_\omega$, so $w \in \mathcal{U}_\omega$, hence $\mathcal{U}_\omega$ is a subspace.

It is easy to see that $\mathcal{U}_\Omega = \bigcap_{\omega \in \Omega} \mathcal{U}_\omega$. Since intersections of subspaces are subspaces, we conclude that $\mathcal{U}_\Omega$ is a subspace. □

Lemma 2. *$\mathcal{U}_\omega$ contains all the r-dimensional subspaces that fit ω, and $\mathcal{U}_\Omega$ contains all the r-dimensional subspaces that fit Ω.*

Proof. Let S be an r-dimensional subspace that fits ω, and $u \in S$. Since $S_\omega = S^\star_\omega$ (see Remark 3), $u_\omega \in S^\star_\omega$, so $u \in \mathcal{U}_\omega$. Since u was arbitrary, we have that $S \subset \mathcal{U}_\omega$. Since S was arbitrary, we conclude that $\mathcal{U}_\omega$ contains all the r-dimensional subspaces that fit ω.

Now let S be an r-dimensional subspace that fits Ω, and $u \in S$. Since $S_{\omega_i} = S^\star_{\omega_i}$ for every $i \in \mathcal{I}$ (see Remark 5), $u_{\omega_i} \in S^\star_{\omega_i}$ for every $i \in \mathcal{I}$, hence $u \in \mathcal{U}_\Omega$. Since u was arbitrary, we have that $S \subset \mathcal{U}_\Omega$. Since S was arbitrary, we conclude that $\mathcal{U}_\Omega$ contains all the r-dimensional subspaces that fit Ω. □

Corollary 1. *$\mathcal{U}_\omega$ fits ω. Furthermore, $\mathcal{U}_\Omega$ fits Ω whenever there is an r-dimensional subspace that fits Ω.*

Proof. $S^\star$ is an r-dimensional subspace that clearly fits ω, so by Lemma 2, $S^\star \subset \mathcal{U}_\omega$, which implies $\mathcal{U}_\omega$ fits ω (see Remark 3).

Now suppose there is an r-dimensional subspace S that fits Ω, such that $S^\star_{\omega_i} = S_{\omega_i}$ for every $i \in \mathcal{I}$ (see Remark 5). By Lemma 2, $S \subset \mathcal{U}_\Omega$, which implies $\mathcal{U}_\Omega$ fits Ω. □

Remark 7. Notice the importance of the requirement that S is an r-dimensional subspace in Lemma 2 and Corollary 1.

In the case when S is assumed to be an r-dimensional subspace that fits $\mathbf{\Omega}$, since $S_{\omega_i} = S^\star_{\omega_i}$ for every i (see Remark 5), we know that every $u \in S$ will satisfy $u_{\omega_i} \in S^\star_{\omega_i}$ for every i, hence $S \subset \mathcal{U}_{\mathbf{\Omega}}$.

If we drop the assumption that S is r-dimensional, we only know that $S^\star_{\omega_i} \subset S_{\omega_i}$ for every i, whence we only know that there exist u's in S that will satisfy $u_{\omega_i} \in S^\star_{\omega_i}$ for every i, but there might also be some other u's in S that won't satisfy this, whence we cannot conclude that $S \subset \mathcal{U}_{\mathbf{\Omega}}$. For example, if $S = \mathbb{R}^d$, it is clear that S will fit $\mathbf{\Omega}$, but won't be contained in $\mathcal{U}_{\mathbf{\Omega}}$. ■

Observe that if there are more than one r-dimensional subspaces that fit Ω, since they are contained in $\mathcal{U}_\Omega$, $\dim \mathcal{U}_\Omega > r$. Moreover, since $\mathcal{U}_\Omega$ is a subspace, if $\dim \mathcal{U}_\Omega > r$, there are infinitely many r-dimensional subspaces contained in $\mathcal{U}_\Omega$. It is easy to see that infinitely many of such subspaces will fit Ω. In other words, using these results, we have the following.

Corollary 2. *There is only one r-dimensional subspace that fits Ω iff $\dim \mathcal{U}_\Omega = r$. Furthermore, if there are more than one r-dimensional subspaces that fit Ω, there are infinitely many.*

7.2 One at a time

We continue our analysis by studying $\mathcal{U}_\omega$. The key idea is that every hyperplane, i.e., every $(d-1)$-dimensional subspace of $\mathbb{R}^d$, is characterized by its orthogonal direction.

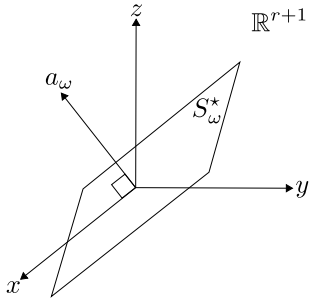

Figure 3: Every hyperplane is characterized by its orthogonal direction.

By definition, $S^\star_\omega$ is a subspace of $\mathbb{R}^{r+1}$. Since $\dim S^\star = r$, and $S^\star$ is non-degenerate, it is easy to see that $\dim S^\star_\omega = r$, i.e., $S^\star_\omega$ is a hyperplane in $\mathbb{R}^{r+1}$. As

such, it is characterized by its orthogonal direction, which can be fully specified by a unique —up to scalar multiplication— non-zero vector of $\mathbb{R}^{r+1}$ orthogonal to $S^\star_\omega$, say a_ω. For such a_ω we have:

$$S^\star_\omega = \{u_\omega \in \mathbb{R}^{r+1} : \langle u_\omega, a_\omega \rangle = 0\} = \ker a_\omega^\mathsf{T}.$$

If we define a as the row vector in $\mathbb{R}^d$ with the entries of a_ω in the positions of ω, and zeros elsewhere, it is clear that

$$\mathcal{U}_\omega = \{u \in \mathbb{R}^d : \langle u_\omega, a_\omega \rangle = 0\} = \{u \in \mathbb{R}^d : \langle u, a \rangle = 0\} = \ker a.$$

Thus $\mathcal{U}_\omega$ is also a hyperplane: the $(d-1)$-dimensional subspace of $\mathbb{R}^d$ characterized by a, containing every $u \in \mathbb{R}^d$ that satisfies $u_\omega \in S^\star_\omega$.

Notice that just as $\mathcal{U}_\omega$, a is a function of $S^\star$ and ω, and since $S^\star$ is non-degenerate,

Lemma 3. *a has exactly $r+1$ non-zero entries in the positions of ω.*

Proof. Suppose for contradiction that a_ω has at least one zero entry. Use $\upsilon \subset \{1, ..., d\}$ to denote the set of the position of the non-zero entries of a_ω, and use it analogous to ω. Since a_ω is orthogonal to $S^\star_\omega$, we have that $\langle a_\omega, u_\omega \rangle = \langle a_\upsilon, u_\upsilon \rangle = 0$ for every $u_\upsilon \in S^\star_\upsilon$, i.e., $S^\star_\upsilon \subset \ker a_\upsilon$. Then

$$\dim S^\star_\upsilon \leq \dim \ker a_\upsilon^\mathsf{T} = |\upsilon| - 1 < |\upsilon|,$$

which is a contradiction, as $S^\star$ is non-degenerate by **A3**. The statement follows directly by the definition of a as the row vector in $\mathbb{R}^d$ with the entries of a_ω in the positions of ω. □

7.3 Several at Once

Our next step is precisely the most obvious one. Let $a_{\omega_i} \in \mathbb{R}^{r+1}$ be a vector in the orthogonal direction of $S^\star_{\omega_i}$, and a_i be the row vector in $\mathbb{R}^d$ with the entries of a_{ω_i} in the positions of ω_i, and zeros elsewhere. With $\mathcal{U}_i$ as shorthand for $\mathcal{U}_{\omega_i}$, then

$$\mathcal{U}_\Omega = \bigcap_{i \in \mathcal{I}} \mathcal{U}_i = \bigcap_{i \in \mathcal{I}} \ker a_i = \ker A,$$

where A is the $n \times d$ matrix with $\{a_i\}_{i \in \mathcal{I}}$ as its rows. In the particular case when $\Omega = \boldsymbol{\Omega}$, we use $\boldsymbol{A}$ instead of A, i.e., $\boldsymbol{A}$ is the $N \times d$ matrix with rows $\{a_i\}_{i=1}^N$. Notice that by Lemma 3, A has exactly m non-zero columns in the positions of $\mathcal{J}$, while $\boldsymbol{A}$ has no zero columns by **A5**.

This way, each a_i defines a hyperplane with all the vectors $u \in \mathbb{R}^d$ that satisfy $u_{\omega_i} \in S^\star_{\omega_i}$, namely $\mathcal{U}_i$. $\mathcal{U}_\Omega$ is the intersection of all such hyperplanes, hence it contains all the vectors $u \in \mathbb{R}^d$ that satisfy $u_{\omega_i} \in S^\star_{\omega_i}$ simultaneously for every i.

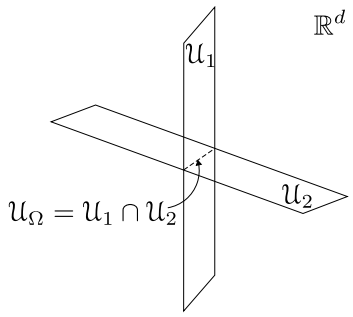

Figure 4: $\mathcal{U}_\Omega = \ker A = \bigcap_{i\in\mathcal{I}} \ker a_i = \bigcap_{i\in\mathcal{I}} \mathcal{U}_i$. This figure could be an illustration of the setup in Example 11.

This way, A characterizes $\mathcal{U}_\Omega$ just as a characterizes $\mathcal{U}_\omega$; and just as $\mathcal{U}_\Omega$, A is a function of $\mathcal{S}^\star$, $\mathcal{K}$ and Ω (see Remark 4).

If we let ℓ be the number of linearly independent rows in A, we immediately know from elemental linear algebra[6] that $\dim \mathcal{U}_\Omega = \dim \ker A = d - \ell$. We state this as a lemma, as it is an observation that will come up in our subsequent analysis.

Lemma 4 ($\dim \mathcal{U}_\Omega$)**.** *$\dim \mathcal{U}_\Omega = d - \ell$.*

It is clear that we are interested on determining when Ω defines an A with linearly independent rows. But before moving to that, let us discuss one very nice property of A.

Lemma 5. *There exists no $\mathcal{A}$ with a row with fewer than $r+1$ non-zero entries such that* $\ker A = \ker \mathcal{A}$.

Proof. Very similar to the proof of Lemma 3, suppose for contradiction that there exists an $\mathcal{A}$ with a row a with fewer than $r+1$ non-zero entries such that $\ker A = \ker \mathcal{A}$. Use $\upsilon \subset \{1, ..., d\}$ to denote the set of the positions of the non-zero entries of a, and use it analogous to ω.

Let S be an r-dimensional subspace that fits Ω. By Lemma 2, $S \subset \mathcal{U}_\Omega$, so $S \subset \mathcal{U}_\Omega = \ker A = \ker \mathcal{A}$, hence every $u \in S$ must satisfy $\langle \mathsf{a}_\upsilon, u_\upsilon \rangle = 0$, i.e., $S_\upsilon \subset \ker \mathsf{a}_\upsilon^\mathsf{T}$. Then

$$\dim S_\upsilon \leq \dim \mathsf{a}_\upsilon^\mathsf{T} = |\upsilon| - 1 < |\upsilon|,$$

i.e., S is degenerate. Since S was arbitrary, we know that this holds for every S that fits Ω, which contradicts **A6**. □

As one simple consequence of Lemma 5 we obtain an extremely useful result:

Corollary 3 (ℓ)**.** *$m \geq \ell + r$.*

Proof. Suppose without loss of generality that A has all its zero columns — if any— in the first block, and that its first ℓ non-zero columns are linearly

independent; otherwise we can simply permute the columns of A accordingly. Let A be the $n \times m$ matrix with only the non-zero columns of A, and suppose for contradiction that $m < \ell + r$. We can then transform A into the following reduced row echelon form:

$$\overbrace{\left[\begin{array}{c|c} 0 & \mathsf{A} \end{array}\right]}^{d-m \quad m} = A \sim \mathcal{A} = \overbrace{\left[\begin{array}{c|c|c} 0 & I & \mathsf{B} \\ \hline & \mathbf{0} & \mathbf{0} \end{array}\right]}^{d-m \qquad m} \begin{array}{l} \left.\begin{array}{l} \\ \\ \end{array}\right\} \ell \\ \left.\begin{array}{l} \\ \end{array}\right\} n-\ell. \end{array}$$

$$\underbrace{\qquad}_{\ell} \qquad \underbrace{\qquad}_{m-\ell<r}$$

We know from elemental linear algebra[6] that $\ker A = \ker \mathcal{A}$, nevertheless the top ℓ rows of $\mathcal{A}$ have fewer than $r+1$ non-zero entries, which is a contradiction by Lemma 5. □

7.4 Independence

From our discussion in §7.3, it is clear that we are interested on determining when $\boldsymbol{\Omega}$ defines an $\boldsymbol{A}$ with linearly independent rows.

With this in mind we use the following.

Definition 12 (Independent set)**.** *We say* Ω *is* independent *iff* all *rows of* A *are linearly independent. Otherwise we say* Ω *is* dependent*.*

With this definition we can now say that our next goal is to be able to identify when Ω is independent. The answer to this is given by the following.

Lemma 6 (Characterization of independent sets)**.** $\breve{\Omega}$ *is independent iff*

$$m \geq n + r \qquad \forall\ \Omega \subset \breve{\Omega}.$$

Remark 8. Notice that Lemma 6 defines a matroid[7]. ■

This lemma represents a central part of Theorems 1 and 2, and is possibly the most transcending result of the document. In order to prove it we will require the following.

Definition 13 (Dependent set, redundant set)**.** *We say* ω *is* dependent *on* Ω *or* redundant *iff a is linearly dependent on the rows of* A*. Otherwise, we say* ω *is independent of* Ω*.*

The next definition will play a crucial role in the proofs of Lemmas 6 and 7 in this section, and also in Lemmas 8, 10 and 11, in §8.

Definition 14 (Basis)**.** *We say* Ω *is a* basis *of* ω *iff* Ω *is an independent set such that* ω *is dependent on* Ω *but is not dependent on any proper subset of* Ω*.*

Remark 9. Observe that bases are not unique. In fact, for a single ω there could be several bases, and even of different sizes. ■

Example 14. Suppose $r = 2$ and

$$\mathbf{\Omega} = \begin{bmatrix} \times & \cdot & \times & \cdot & \cdot & \times \\ \times & \times & \cdot & \cdot & \cdot & \cdot \\ \times & \times & \times & \times & \cdot & \cdot \\ \cdot & \times & \times & \times & \times & \cdot \\ \cdot & \cdot & \cdot & \times & \times & \times \\ \cdot & \cdot & \cdot & \cdot & \times & \times \end{bmatrix}.$$

ω_3 is trivially dependent on itself, so $\{\omega_3\}$ is a trivial basis of ω_3. One can verify using Lemma 6 that ω_3 is dependent on $\{\omega_1, \omega_2\}$. Then $\{\omega_1, \omega_2\}$ is also a basis of ω_3. Finally, ω_3 is also dependent on $\{\omega_4, \omega_5, \omega_6\}$, so $\{\omega_4, \omega_5, \omega_6\}$ is also a basis of ω_3. Notice that these three different bases are not even of the same size. ■

The crux of Lemma 6, and all our further results for that matter, lies in the following statement.

Lemma 7 (Property of bases)**.** *Let* Ω *be a basis of* ω*. Then* $m = n + r$*.*

Proof. Let Ω be a basis of ω. If Ω is a trivial basis, i.e., if $\Omega = \{\omega\}$, it is trivially true that $m = n + r$. Suppose then that Ω is non-trivial.

First observe that $\omega \subset \mathcal{J}$; otherwise a would have a nonzero entry corresponding to a zero column of A, whence a could not possibly be linearly dependent on A.

We can assume without loss of generality that A has all its zero columns —if any— in the first block and a has its non-zero entries in the last $r + 1$ columns; otherwise, we may just permute the columns of a and A accordingly. Also assume without loss of generality that the first non-zero entry of a is 1 —otherwise we can just scale the row— and let $\hat{\mathsf{a}}$ denote the $1 \times r$ row with the remaining non-zero entries of a, such that we can write:

$$\left[\begin{array}{c} A \\ \hline a \end{array}\right] = \left[\begin{array}{c|c|c} \overbrace{0}^{d-m} & \overbrace{\begin{array}{c} B \\ \hline \underbrace{\mathbf{0}}_{m-r} \end{array}}^{m} & \begin{array}{c} C \\ \hline \underbrace{\hat{\mathsf{a}}}_{r} \end{array} \end{array}\right] \begin{array}{l} \left.\vphantom{\begin{array}{c} B \\ \hline \mathbf{0} \end{array}}\right\} n \\ \left.\vphantom{1}\right\} 1. \end{array} \tag{3}$$

Recall that we want to show that $m = n + r$. Since Ω is independent, we know $n = \ell$, so by Corollary 3 we immediately know that $m \geq n + r$. As we specified in §6, all our statements hold for almost every $\mathcal{S}^\star$, so it suffices to show that $m > n + r$ only in a set of measure zero. Suppose then that $m > n + r$. This implies that B has strictly more than n columns.

Observe that the rows of B are linearly independent. To see this, very similar to what we did in Corollary 3, suppose for contradiction that they are not. This implies that we can transform A into the following reduced row echelon form:

$$A \sim \mathcal{A} = \left[\begin{array}{c|c|c} \overbrace{0}^{d-m} & \overbrace{\mathcal{B}}^{m} & \mathsf{C} \\ \hline & \mathbf{0} & \end{array}\right] \Big\} > 0.$$
$$\underbrace{}_{\ell} \quad \underbrace{\phantom{\mathsf{C}}}_{r}$$

We know from elemental linear algebra[6] that $\ker A = \ker \mathcal{A}$. Nevertheless, the last row of $\mathcal{A}$ has at most r non-zero entries, which is a contradiction by Lemma 5.

Let a be the $1 \times m$ row with only the entries of A in the positions of $\mathcal{J}$, and A be the $n \times m$ matrix with only the columns of A in the positions of $\mathcal{J}$.

Going back to (3), since the rows of B are linearly independent and $m - r > n$, we know, B has n linearly independent columns. Let B denote the $n \times n$ block of B that contains n linearly independent columns, and $\mathcal{B}$ the $n \times (m - n - r)$ remaining block of B. We can thus assume without loss of generality that:

$$\left[\begin{array}{c} \mathsf{A} \\ \hline \mathsf{a} \end{array}\right] = \left[\begin{array}{c|c|c} \overbrace{\mathcal{B}}^{B} & \mathsf{B} & C \\ \hline \mathbf{0} & \mathbf{0} \;\; 1 & \hat{\mathsf{a}} \end{array}\right] \begin{array}{l} \Big\} n \\ \Big\} 1. \end{array}$$
$$\underbrace{\phantom{\mathbf{0}}}_{m-n-r \geq 1} \quad \underbrace{\phantom{\mathbf{0}\;\;1}}_{n} \quad \underbrace{\phantom{\hat{\mathsf{a}}}}_{r}$$

Notice that the column of B corresponding to the 1 in a must belong to B (otherwise, we have that $\beta\mathsf{B} = 0$, with β as in (5), which implies that B has a linearly dependent row, hence a linearly dependent column).

We can further assume without loss of generality that the first non-zero entry of every row of A is 1; otherwise we may just scale each row. Finally, we may also assume that the first column of A, namely the first column of $\mathcal{B}$ has all its $\flat$ non-zero entries —all ones by construction— on the top rows; otherwise we

may just permute the rows of A accordingly, such that we have:

$$\left[\begin{array}{c}\mathsf{A}\\\hline \mathsf{a}\end{array}\right] = \overset{\mathfrak{b}}{\left\{\vphantom{\begin{array}{c}a\\b\end{array}}\right.} \left[\begin{array}{c|c|c|c} \overbrace{\begin{array}{c}\mathbf{1}\\\mathbf{0}\end{array}}^{\mathcal{B}} & \mathcal{B} & \mathsf{B} & C \\\hline \underbrace{\mathbf{0}}_{1} & \underbrace{\mathbf{0}}_{m-n-r-1\geq 0} & \underbrace{\mathbf{0}\;\big|\;1}_{n} & \underbrace{\hat{\mathsf{a}}}_{r} \end{array}\right] \left.\vphantom{\begin{array}{c}a\\b\end{array}}\right\} \begin{array}{l} n \\ 1. \end{array} \tag{4}$$

Since a is minimally linearly dependent on A, we know there exists a unique row vector $\beta \in \mathbb{R}^n$ with all non-zero entries such that

$$\beta A = a, \tag{5}$$

which implies $1 < \mathfrak{b} \leq n$. In particular, using (5) on the B block of (4), we have that $\beta\mathsf{B} = \left[\begin{array}{c|c} \mathbf{0} & 1 \end{array}\right]$, and since B is full-rank, we can solve for β:

$$\beta = \left[\begin{array}{c|c} \mathbf{0} & 1 \end{array}\right] \mathsf{B}^{-1}, \tag{6}$$

i.e., β is the the last row of B^{-1}. We know from elemental linear algebra[6] that

$$\mathsf{B}^{-1} = \frac{1}{|\mathsf{B}|}\mathsf{B}^*,$$

where $|\cdot|$ and $\cdot^*$ denote the determinant and the adjoint matrix of $\cdot$, respectively. Since B^* is the transpose of the cofactor matrix of B, we have that

$$\mathsf{b}^{-1}_{ni} = \pm_{in}\frac{|\mathsf{B}_{in}|}{|\mathsf{B}|}, \tag{7}$$

where B_{in} denotes the $(n-1)\times(n-1)$ minor of B obtained by removing the i^{th} row and the n^{th} column of B, and $\pm_{in}$ denotes the $(i,n)^{th}$ entry of the following matrix:

$$\pm := \begin{bmatrix} + & - & + & - & \cdots \\ - & + & - & + & \cdots \\ + & - & + & - & \cdots \\ - & + & - & + & \cdots \\ \vdots & \vdots & \vdots & \vdots & \ddots \end{bmatrix},$$

For example, if:

$$\mathsf{B} = \left[\begin{array}{ccc|c} \mathsf{b}_{11} & \mathsf{b}_{12} & \mathsf{b}_{13} & \mathsf{b}_{14} \\\hline \mathsf{b}_{21} & \mathsf{b}_{22} & \mathsf{b}_{23} & \mathsf{b}_{24} \\ \mathsf{b}_{31} & \mathsf{b}_{32} & \mathsf{b}_{33} & \mathsf{b}_{34} \\ \mathsf{b}_{41} & \mathsf{b}_{42} & \mathsf{b}_{43} & \mathsf{b}_{44} \end{array}\right],$$

then

$$\pm_{1n}|\mathsf{B}_{1n}| = -\begin{vmatrix} \mathsf{b}_{21} & \mathsf{b}_{22} & \mathsf{b}_{23} \\ \mathsf{b}_{31} & \mathsf{b}_{32} & \mathsf{b}_{33} \\ \mathsf{b}_{41} & \mathsf{b}_{42} & \mathsf{b}_{43} \end{vmatrix}.$$

Now observe that for $i \neq 1$, we have:

$$|\mathsf{B}_{in}| = \sum_{\ddot{\imath}=1}^{n-1} \pm_{1\ddot{\imath}} \mathsf{b}_{1\ddot{\imath}} |\mathsf{B}_{in}^{1\ddot{\imath}}|,$$

where $\mathsf{B}_{in}^{1\ddot{\imath}}$ is the $(n-2)\times(n-2)$ minor of B_{in} obtained by removing the first row and the $\ddot{\imath}^{th}$ column of B_{in}. In our example,

$$\mathsf{B} = \left[\begin{array}{cc|c|c} \mathsf{b}_{11} & \mathsf{b}_{12} & \mathsf{b}_{13} & \mathsf{b}_{14} \\ \hline \mathsf{b}_{21} & \mathsf{b}_{22} & \mathsf{b}_{23} & \mathsf{b}_{24} \\ \hline \multicolumn{2}{c|}{\multirow{2}{*}{B_{2n}^{13}}} & \mathsf{b}_{33} & \mathsf{b}_{34} \\ \multicolumn{2}{c|}{} & \mathsf{b}_{33} & \mathsf{b}_{44} \end{array}\right].$$

Recall that $\beta_i \neq 0$ for every i, and since $\beta_i = \frac{\pm_{in}|\mathsf{B}_{in}|}{|\mathsf{B}|}$, it is clear that $|\mathsf{B}_{in}| \neq 0$ for every i. This implies that there is at least one $\ddot{\imath}$ for which $\mathsf{b}_{1\ddot{\imath}} \neq 0$.

Now let us look back at (5). Using the first column of (4) we obtain $\beta \left[\begin{array}{c|c} \mathbf{1} & \mathbf{0} \end{array}\right]^{\mathsf{T}} = 0$. Substituting (6), (7), and factoring out the common term $|\mathsf{B}|$, we obtain:

$$\sum_{i=1}^{\flat} \pm_{in}|\mathsf{B}_{in}| = 0. \tag{8}$$

We will now show that $f := \sum_{i=1}^{\flat} \pm_{in}|\mathsf{B}_{in}|$ is a non-zero polynomial. Write:

$$\begin{aligned} f := \sum_{i=1}^{\flat} \pm_{in}|\mathsf{B}_{in}| &= \pm_{1n}|\mathsf{B}_{1n}| + \sum_{i=2}^{\flat} \pm_{in}|\mathsf{B}_{in}| \\ &= \pm_{1n}|\mathsf{B}_{1n}| + \sum_{i=2}^{\flat} \pm_{in} \sum_{\ddot{\imath}=1}^{n-1} \pm_{1\ddot{\imath}} \mathsf{b}_{1\ddot{\imath}} |\mathsf{B}_{in}^{1\ddot{\imath}}| \\ &= \pm_{1n}|\mathsf{B}_{1n}| + \sum_{\ddot{\imath}=1}^{n-1} \mathsf{b}_{1\ddot{\imath}} \underbrace{\sum_{i=2}^{\flat} \pm_{in} \pm_{1\ddot{\imath}} |\mathsf{B}_{in}^{1\ddot{\imath}}|}_{=:\mathsf{c}_{\ddot{\imath}}}, \end{aligned}$$

where $\mathsf{c}_{\ddot{\imath}}$ is a constant polynomial of $\mathsf{b}_{1\ddot{\imath}}$. By simple inspection, one can see that $|\mathsf{B}_{1n}|$ does not depend on $\mathsf{b}_{1\ddot{\imath}}$, i.e., it is also a constant polynomial of $\mathsf{b}_{1\ddot{\imath}}$. Thus, if $\mathsf{c}_{\ddot{\imath}} \neq 0$ for some $\ddot{\imath}$ for which $\mathsf{b}_{1\ddot{\imath}} \neq 0$, we can immediately conclude that f is a non-zero polynomial of $\mathsf{b}_{1\ddot{\imath}}$.

On the other hand, if $\mathsf{c}_{\ddot{\imath}} = 0$ for every $\ddot{\imath}$ for which $\mathsf{b}_{1\ddot{\imath}} \neq 0$, then $f = \pm_{1n}|\mathsf{B}_{1n}|$. Similar to what we did for $i \neq 1$, we can write:

$$|\mathsf{B}_{1n}| = \sum_{\ddot{\imath}=1}^{n-1} \pm_{2\ddot{\imath}} \mathsf{b}_{2\ddot{\imath}} |\mathsf{B}_{1n}^{2\ddot{\imath}}|,$$

where $\mathsf{B}_{in}^{2\ddot{\imath}}$ is the $(n-2)\times(n-2)$ minor of B_{1n} obtained by removing the second row and the $\ddot{\imath}^{th}$ column of B_{1n}. Again, since $\beta_1 \neq 0$, there is at least one $\ddot{\imath}$ for which $\mathsf{b}_{2\ddot{\imath}} \neq 0$, thus f is a non-zero polynomial of $\mathsf{b}_{2\ddot{\imath}}$.

Observe that every $\mathsf{b}_{i\ddot{\imath}}$ is either zero, or one of the entries of a_{ω_i}. Thus f is a non-zero polynomial of at least one of the entries of a_{ω_1} or a_{ω_2}, say $a_{\omega_{i\ddot{\imath}}}$.

Since $\flat > 1$, $a_{\omega_{i\ddot{\imath}}} \neq 1$, hence f is a non-zero polynomial in $\mathbb{R}[a_{\omega_{i\ddot{\imath}}}]$. We thus know from elemental measure theory and algebraic geometry[8] that the variety defined by (8) has measure zero. In other words, there is only a subset of measure zero over $a_{\omega_{i\ddot{\imath}}}$, hence over the set of all sets $\mathcal{S}^\star$, for which $m > n + r$, which concludes the proof. □

We are now ready to present the proof of Lemma 6.

Proof. (Lemma 6). We prove both directions by contrapositive. Explicitly, we will show that $\breve{\Omega}$ is dependent iff $\exists\ \Omega \subset \breve{\Omega}$ with $m < n + r$.

($\Rightarrow$) Let $\breve{\Omega}$ be dependent. By definition, it contains a redundant ω with basis $\bar{\Omega} \subset \mathbf{\Omega}\backslash\omega$. By Lemma 7, $\bar{n} = \bar{m} - r$.

Take $\Omega := \bar{\Omega} \cup \omega$. It is clear that $m = \bar{m}$. Nevertheless, $n = \bar{n} + 1$. Thus $m < n + r$, and we have the first implication.

($\Leftarrow$) Suppose $\exists\ \Omega \subset \breve{\Omega}$ with $m < n + r$. By Corollary 3, $n > \ell$, which implies Ω is be dependent. Of course, since $\Omega \subset \breve{\Omega}$, $\breve{\Omega}$ is also dependent, which concludes the second part of the proof. □

Example 15. By Lemma 6, the $\mathbf{\Omega}$'s from Examples 6, 11 and **??** are independent, for their respective r's. On the other hand, the $\mathbf{\Omega}$'s from Examples 12, 14, 22 and 18 are dependent, for their respective r's.

Notice that $\mathbf{\Omega}$ may be dependent and still satisfy that there is only one subspace that fits it. Such is the case of Examples 12, 14 and 18: they are dependent, but they contain an independent set of size $d - r$. In other words, they have some redundant vectors. ■

Remark 10. As we will see in §8.1, redundant vectors are not useless; they are required to guarantee that different incomplete vectors belong to the same subspace. In fact, it is easy to see that the conditions of Theorem 1 require the existence of redundant vectors in $\mathbf{\Omega}$. ■

7.5 All you need is $d - r$

Recall that we are interested on determining conditions to guarantee that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$. Observe that this is not implied by $\mathbf{\Omega}$ being independent nor viceversa. Nevertheless, as a simple consequence of Corollary 2, Lemma 4 and Definition 12, we obtain the following.

Corollary 4. *Suppose $n = d - r$. There is only one r-dimensional subspace that fits Ω iff Ω is independent.*

In other words, Corollary 4, is telling us that all we need is that $\mathbf{\Omega}$ has $d-r$ independent ω's to guarantee that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$. Theorem 2 is precisely a combination of this observation and the characterization of independent sets.

Example 16. Consider $\mathbf{\Omega}$ as in Example 6. $\mathbf{\Omega}$ satisfies the conditions of Corollary 4, so there is only one r-dimensional subspace that fits $\mathbf{\Omega}$. ■

7.6 Proof of Theorem 2

We conclude this section with the proof of Theorem 2, which comes immediately as a consequence of Corollaries 2 and 3 and Lemmas 4 and 6. Observe that Theorem 2 essentially states that there exits only one r-dimensional subspace that fits $\mathbf{\Omega}$ iff $\mathbf{\Omega}$ contains an independent set of size $d - r$.

Proof. (Theorem 2) ($\Rightarrow$) We prove this by contrapositive. Suppose $\nexists\ \breve{\Omega} \subset \mathbf{\Omega}$ of size $d-r$ such that $m \geq n+r$ for every $\Omega \subset \breve{\Omega}$. By Lemma 6, we know there is no independent set of size $d-r$ in $\mathbf{\Omega}$, i.e., there are no $d-r$ linearly independent rows in $\boldsymbol{A}$. By Lemma 4, $\dim \mathcal{U}_\Omega > r$, thus by Corollary 2 there exist infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$.

($\Leftarrow$) Suppose there exists an $\breve{\Omega} \subset \mathbf{\Omega}$ of size $d - r$ such that $m \geq n + r$ for every $\Omega \subset \breve{\Omega}$. By Lemma 6 $\breve{\Omega}$ is independent, i.e., $\boldsymbol{A}$ contains at least $d - r$ linearly independent rows. Furthermore, by Corollary 3 $\boldsymbol{A}$ contains exactly $d-r$ linearly independent rows. By Lemma 4, $\dim \mathcal{U}_\Omega = r$, thus by Corollary 2 there is only one r-dimensional subspace that fits $\breve{\Omega}$. Since $\breve{\Omega} \subset \mathbf{\Omega}$, there is also only one r-dimensional subspace that fits $\mathbf{\Omega}$. □

8 All of a kind

Notice that so far, Theorem 2 is only stating when there will be only one r-dimensional subspace that fits $\mathbf{\Omega}$. It is not yet implying anything about $\mathcal{K}$. For all we know such subspace could fit $\mathbf{\Omega}$, even if the x's belong to different subspaces.

In other words, even if there is only one r-dimensional subspace that fits $\mathbf{\Omega}$, there is yet nothing that assures us that all the columns of $\mathcal{X}$ indeed lie in one r-dimensional subspace.

Example 17. Consider the same setup as in Example 11. One can easily verify using Theorem 2 that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$. But suppose that the columns of $\mathcal{X}$ don't lie in the same r-dimensional subspace, i.e., $k_1 \neq k_2$, as in Example 1.

One can see that there are infinitely many r-dimensional subspaces that will satisfy $S_{\omega_1} = S^\star_{\omega_1}$; the r-dimensional subspaces contained in $\mathcal{U}_1$, e.g., any of the subspaces in Figure 5(a). These will fit ω_1. There are also infinitely many r-dimensional subspaces that will satisfy $S_{\omega_2} = S^\star_{\omega_2}$; the r-dimensional subspaces contained in $\mathcal{U}_2$, e.g., any of the subspaces in Figure 5(b). These will fit ω_2.

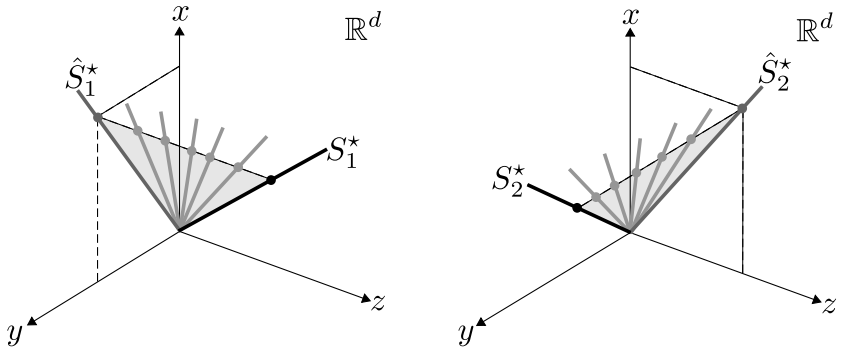

(a) r-dimensional subspaces that satisfy $S_{\omega_1} = S^\star_{\omega_1}$, hence fit ω_1.

(b) r-dimensional subspaces that satisfy $S_{\omega_2} = S^\star_{\omega_2}$, hence fit ω_2.

Figure 5: There are infinitely many r-dimensional subspaces that satisfy $S_{\omega_i} = S^\star_{\omega_i}$, hence fit ω_i; namely, the subspaces whose projection onto $\mathbb{R}^d_{\omega_i}$ is the same as the projection of $S^\star_i$ onto $\mathbb{R}^d_{\omega_i}$.

And despite the columns of $\mathcal{X}$ don't lie in the same r-dimensional subspace, there is one r-dimensional subspace, S, that satisfies $S_{\omega_1} = S^\star_{\omega_1}$ and $S_{\omega_2} = S^\star_{\omega_2}$ simultaneously, and thus fits $\mathbf{\Omega}$: $S = \mathcal{U}_{\mathbf{\Omega}} = \mathcal{U}_1 \cap \mathcal{U}_2$.

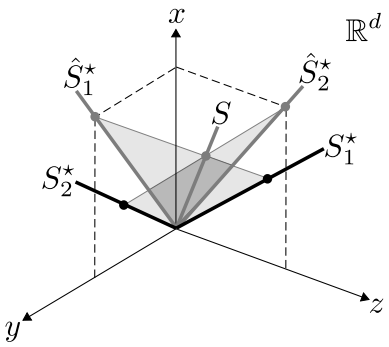

Figure 6: Despite $k_1 \neq k_2$, there is one r-dimensional subspace S that satisfies $S_{\omega_1} = S^\star_{\omega_1}$ and $S_{\omega_2} = S^\star_{\omega_2}$ simultaneously, and thus fits $\mathbf{\Omega}$: $S = \mathcal{U}_{\mathbf{\Omega}} = \mathcal{U}_1 \cap \mathcal{U}_2$.

For an explicit instance of the example above, take Example 2. It is easy to see that S fits $\mathbf{\Omega}$ despite the columns of $\mathcal{X}$ do not lie in a 1-dimensional subspace. ■

We want to know when can we be sure that if there is only one r-dimensional subspace that fits $\mathbf{\Omega}$, it is because all the columns of $\mathcal{X}$ indeed lie in an r-dimensional subspace.

Example 18. Continuing with Example 17, suppose we had an additional $\omega_3 = \{1, 3\}$, i.e., suppose

$$\mathbf{\Omega} = \begin{bmatrix} \times & \times & \cdot \\ \times & \cdot & \times \\ \cdot & \times & \times \end{bmatrix},$$

Then almost surely, $\mathcal{U}_{\mathbf{\Omega}} = \mathcal{U}_1 \cap \mathcal{U}_2 \cap \mathcal{U}_3$ will only contain an r-dimensional subspace iff all the columns of $\mathcal{X}$ lie in an r-dimensional subspace.

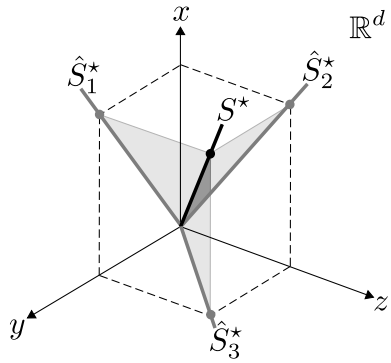

Figure 7: Almost surely, $\mathcal{U}_{\mathbf{\Omega}}$, the intersection of three hyperplanes in $\mathbb{R}^3$, each corresponding to one of the $\mathcal{U}_i$'s, is a line iff $k_1 = k_2 = k_3$, whence $\mathcal{U}_{\mathbf{\Omega}} = S^\star$.

■

In other words, we want to derive necessary and sufficient conditions on $\mathbf{\Omega}$ to guarantee that all the columns of $\mathcal{X}$ indeed lie in an r-dimensional subspace, i.e., that $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$.

We want those conditions to be sufficient in the sense that if $\mathbf{\Omega}$ satisfies such conditions, then it *must* be true that $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$.

We want those conditions to be necessary in the sense that if $\mathbf{\Omega}$ does not satisfy these conditions, it cannot be guaranteed that $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$, i.e., k_i *could* be different from $k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$, implying that the columns of $\mathcal{X}$ *might* not all lie in the same r-dimensional subspace.

The answer to this question, the main one of this document, is given by Theorem 1, which we will show in this section.

8.1 Characterization of $\mathcal{K}$

We start the work towards the proof of Theorem 1 with the following lemma. It is essential for our further analysis, as will allow us to determine when columns of $\mathcal{X}$ lie in the same r-dimensional subspace.

Lemma 8 (Characterization of $\mathcal{K}$)**.** *Let Ω be a basis of ω_i. Then $k_i = k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$.*

Proof. Let Ω be a basis of ω_i. Since a_i is minimally linearly dependent on A by definition, we may write

$$a_i = \sum_{\ddot{\imath} \in \mathcal{I}} \beta_{\ddot{\imath}} a_{\ddot{\imath}},$$

where $\beta_{\ddot{\imath}} \neq 0$ for every $\ddot{\imath} \in \mathcal{I}$. On the other hand, a_i is a non-zero function of ω_i and $S_i^\star$, say $f_{\dot{\imath}}(\omega_i, S_i^\star)$. Similarly, $a_{\ddot{\imath}}$ is a non-zero function of $\omega_{\ddot{\imath}}$ and $S_{\ddot{\imath}}^\star$, say $f_{\ddot{\imath}}(\omega_{\ddot{\imath}}, S_{\ddot{\imath}}^\star)$. We thus have that

$$f_{\dot{\imath}}(\omega_i, S_i^\star) = \sum_{\ddot{\imath} \in \mathcal{I}} \beta_{\ddot{\imath}} f_{\ddot{\imath}}(\omega_{\ddot{\imath}}, S_{\ddot{\imath}}^\star).$$

The subspaces in $\mathcal{S}^\star$ keep no relation between each other for almost every $\mathcal{S}^\star$. Thus, almost surely, the only way that these equality can hold is iff $S_i^\star = S_{\ddot{\imath}}^\star$ for every $(i, \ddot{\imath})$, i.e., iff $k_i = k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$. □

Of comparable importance is the converse of the previous lemma. It will allow us to determine when columns of $\mathcal{X}$ *might* lie in different subspaces. This, together with our former result give us a complete way to characterize $\mathcal{K}$.

Lemma 9 (Converse characterization of $\mathcal{K}$)**.** *Let ω_i be independent of Ω. Then k_i* might *be different from $k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$.*

Proof. Since we only need to show that k_i *might* be different from $k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$, it suffices an example. Take Example 17. □

8.2 All you need is $d - r$ of a kind

We now know from Lemma 8 that the columns of $\mathcal{X}$ corresponding to a set and its basis belong to the same subspace. We also know from Theorem 2 that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$ if $\mathbf{\Omega}$ contains an independent set of size $d - r$. Combining these two ideas we obtain the following lemma, which intuitively tells us that if we find a set with a basis of size $d-r$, then *all* columns correspond to the same subspace. Conversely, it tells us that if there is one set for which we cannot find a basis of size $d - r$, the columns of $\mathcal{X}$ *might* belong to different subspaces.

Lemma 10 (All of a kind)**.** *$k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$ if $\exists\ \omega$ with a basis of size $d-r$. Conversely, k_i* might *be different from $k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$ if $\nexists$ basis of size $d - r$ for some ω.*

Proof. ($\Leftarrow$) Suppose $\exists\ \omega$ with a basis Ω of size $d - r$. By Lemma 8, $k_i = k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$. Then $S_i^\star$ clearly fits Ω. Since Ω is independent and $n = d - r$, by Corollary 4 there is only one r-dimensional subspace that fits Ω, so $S_i^\star$ is the only r-dimensional subspace that fits Ω. Finally, observe that $S_i^\star$ cannot fit every Ω unless $k_i = k_{\ddot{\imath}}$ for every $\ddot{\imath} \in \mathcal{I}$.

($\Rightarrow$) Suppose $\nexists\ \omega$ with basis of size $d - r$.

Let Ω be the set of sets that ω can be dependent on, i.e.

$$\Omega = \{\breve{\omega} \in \breve{\Omega} : \breve{\Omega} \subset \mathbf{\Omega} \text{ is a basis of } \omega\}.$$

We will show that k_i *might* be different for the elements of $\mathcal{I}$ and the elements of $\mathcal{I}^\mathsf{c}$. By Lemma 9, it suffices to show that every $\breve{\omega} \in \Omega$ is independent of Ω^c.

Let $\breve{\omega} \in \Omega$, and $\breve{\Omega}$ be a basis of ω that contains $\breve{\omega}$. Suppose for contradiction that $\breve{\omega}$ is dependent on Ω^c and let $\bar{\Omega} \subset \Omega^\mathsf{c}$ be a basis of $\breve{\omega}$.

By our definition of basis, ω is not dependent on $\breve{\Omega} \backslash \breve{\omega}$. Nevertheless, since $\bar{\Omega}$ is a basis of $\breve{\omega}$, ω is dependent on $(\breve{\Omega} \backslash \breve{\omega}) \cup \bar{\Omega}$. This implies that there is a basis of ω in $(\breve{\Omega} \backslash \breve{\omega}) \cup \bar{\Omega}$ that contains *at least* one set of $\bar{\Omega}$, say $\bar{\omega}$. Then $\bar{\omega}$ belongs to Ω. Since $\bar{\Omega} \subset \Omega^\mathsf{c}$, $\bar{\omega}$ belongs to Ω^c as well, which is a contradiction.

This implies that $\breve{\omega}$ is independent of Ω^c. Since $\breve{\omega}$ was arbitrary, we conclude that every $\breve{\omega} \in \Omega$ is independent of Ω^c. Thus, by Lemma 9, k_i *might* be different from $k_{\ddot{\imath}}$ for $i \in \mathcal{I}$ and $\ddot{\imath} \in \mathcal{I}^\mathsf{c}$. □

Remark 11. This lemma implies that *if* there is an ω with a basis of size $d-r$, then *every* ω has a basis of size $d-r$. Conversely, *if* there is an ω with no basis of size $d-r$, then *no* ω has a basis of size $d-r$. In other words, we have the following. ■

Corollary 5 (One, then All)**.** *There exists an ω with a basis of size $d-r$ iff every ω has a basis of size $d-r$.*

8.3 Bases Characterization

Lemma 10 tells us that the columns of $\mathcal{X}$ *must* belong to the same subspace iff there is an ω with a basis of size $d-r$. The only remaining step towards the proof of Theorem 1 is to determine when will $\mathbf{\Omega}$ contain an ω with such a basis. The following lemma makes use of Lemma 6 to give us a characterization of bases. This is then used in Theorem 1 to determine when will $\mathbf{\Omega}$ contain an ω with a basis of size $d-r$.

Lemma 11 (Bases Characterization)**.** *Let ω be given and Ω be an independent set. Let $\bar{\Omega} = \Omega \cup \omega$. Ω is a basis of ω iff $\bar{m} < \bar{n} + r$ and $\breve{m} \geq \breve{n} + r$ for every $\breve{\Omega} \subsetneq \bar{\Omega}$.*

Proof. ($\Rightarrow$) Suppose Ω is a basis of ω. Then $\bar{\Omega} = \Omega \cup \omega$ is dependent. By Lemma 6, $\bar{m} < \bar{n} + r$. Let $\breve{\Omega} \subsetneq \bar{\Omega}$. By our definition of basis, $\breve{\Omega}$ is independent. Again, by Lemma 6, $\breve{m} \geq \breve{n} + r$. Since $\breve{\Omega}$ was arbitrary, we conclude that $\breve{m} \geq \breve{n} + r$ for every $\breve{\Omega} \subsetneq \bar{\Omega}$, as desired.

($\Leftarrow$) Assume that $\breve{m} \geq \breve{n} + r$ for every $\breve{\Omega} \subsetneq \bar{\Omega} = \Omega \cup \omega$. This implies by Lemma 6 that Ω is independent and ω is not dependent on any subset of Ω. Further assume that $\bar{m} < \bar{n} + r$. By Lemma 6, $\bar{\Omega}$ is dependent. Then ω is dependent on Ω, and so Ω is a basis of ω. □

8.4 Proof of Theorem 1

We are finally ready to give the proof of Theorem 1, which comes directly as a consequence of Lemmas 6, 10 and 11. Notice that, in a nutshell, the condition of Theorem 1 is that $\mathbf{\Omega}$ contains a set ω with a basis $\breve{\Omega}$ of size $d-r$.

Proof. ($\Rightarrow$) Assume $\exists\ \breve{\Omega} \subset \mathbf{\Omega}$ of size $d-r+1$ such that $m \geq n + r$ for every $\Omega \subsetneq \breve{\Omega}$. Let $\omega \in \breve{\Omega}$. By Lemma 11, $\breve{\Omega} \backslash \omega$ is a basis of ω of size $d-r$. By Lemma 10, $k_i = k_{\ddot{\imath}}$ for every $(i, \ddot{\imath})$. By the same arguments as in the proof of such lemma, $S_i^\star$ is the only r-dimensional subspace that fits $\mathbf{\Omega}$.

($\Leftarrow$) Assume $\nexists\ \omega$ with a basis of size $d - r$. By Lemma 10 k_i *might* be different from $k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$. Furthermore, $\mathbf{\Omega}$ *might* not even contain an independent set of size $d - r$, whence, by Theorem 2, there could be infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$. Moreover, even if $\mathbf{\Omega}$ contains such a set, implying by Lemma 6 that there is only one r-dimensional subspace that fits $\mathbf{\Omega}$, since k_i *might* be different from $k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$, such subspace *might* not even be equal to any of the subspaces in $\mathcal{S}^\star$. □

We conclude the section with a nice converse that comes as a direct consequence of Theorem 1.

Corollary 6 (Converse of Theorem 1)**.** *Assume $\mathbf{\Omega}$ satisfies the assumptions of Theorem 1. There exists no r-dimensional that fits $\mathbf{\Omega}$ iff $k_i \neq k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$.*

9 Intuitively speaking

In this section we give some intuitive explanations of our results from §7, as well as the key ideas behind them.

The idea behind $\mathcal{U}_\omega$ is that every r-dimensional subspace S that fits ω must satisfy $S_\omega = S^\star_\omega$. Recall that $\hat{S}$ is the projection of S onto $\mathbb{R}^d_\omega$, so essentially, the condition $S_\omega = S^\star_\omega$ is telling us that the projections of S and $S^\star$ onto $\mathbb{R}^d_\omega$ are the same, i.e., the projection onto $\mathbb{R}^d_\omega$ of every vector in S lies in $\hat{S}^\star$.

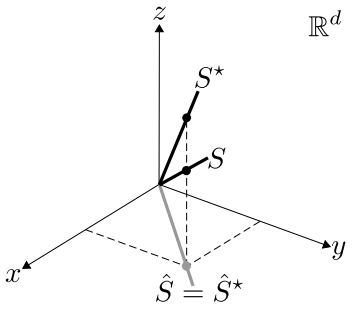

Figure 8: The condition $S_\omega = S^\star_\omega$ is equivalent to saying that $\hat{S}$ and $\hat{S}^\star$, the projections of S and $S^\star$ onto $\mathbb{R}^d_\omega$, are the same.

Essentially, $\mathcal{U}_\omega$ characterizes all the subspaces that fit ω by characterizing all the vectors whose projection onto $\mathbb{R}^d_\omega$ lies in $\hat{S}^\star$; these are precisely the vectors that satisfy $u_\omega \in S^\star_\omega$.

Let us recall that the main goal of §7 is to determine when there is only one r-dimensional subspace that fits Ω. Since all the r-dimensional subspaces that fit Ω are contained in $\mathcal{U}_\Omega$, which is just $\bigcap_{i\in\mathcal{I}} \mathcal{U}_i$, iff we can make sure that $\mathcal{U}_\Omega$ is an r-dimensional subspace, we will be sure that there is only one r-dimensional subspace that fits Ω. In order to do so, we can characterize all the vectors u that lie in $\mathcal{U}_\Omega$.

The key intuitive idea to do so is that the entries of a vector of an r-dimensional subspace are determined given only r of its entries. Since every $u \in \mathcal{U}_\omega$ must fit in $S^\star$, every u_ω must lie in $S^\star_\omega$. What that means is that one entry of every u_ω is constrained as a function of $S^\star$ and the other r entries of u_ω. More specifically for every $u \in \mathcal{U}_\omega$ and for any $j \in \omega$ and $j^\mathsf{c} := \omega \backslash j$, u_j is determined given u_{j^c} and $S^\star$.

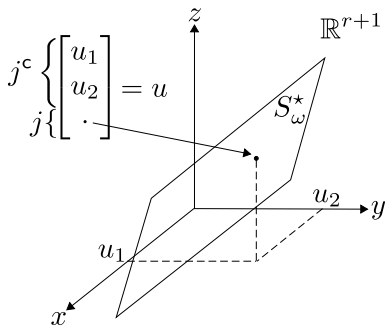

Figure 9: Suppose $r = 2$. One entry of every $u_\omega \in \mathcal{U}_\omega$, say u_j, is determined given the remaining two entries of u_ω, say $u_{j^{\mathsf{c}}}$.

To see how exactly an entry of u_ω is constrained by the other r entries of u_ω, observe that since u_ω must lie in $S^\star_\omega$, u must satisfy

$$u_\omega = U^\star_\omega \gamma$$

for some $\gamma \in \mathbb{R}^r$. Of course, we can rewrite this as

$$\begin{bmatrix} u_{j^{\mathsf{c}}} \\ u_j \end{bmatrix} = \begin{bmatrix} U^\star_{j^{\mathsf{c}}} \\ U^\star_j \end{bmatrix} \gamma.$$

We can focus on the top block and solve for γ:

$$\gamma = (U^\star_{j^{\mathsf{c}}})^{-1} u_{j^{\mathsf{c}}},$$

where we know $(U^\star_{j^{\mathsf{c}}})^{-1}$ exists, as $S^\star$ is non-degenerate. We then focus on the bottom block to obtain:

$$u_j = \underbrace{U^\star_j (U^\star_{j^{\mathsf{c}}})^{-1}}_{\hat{\mathsf{a}}} u_{j^{\mathsf{c}}}, \tag{9}$$

where the only unknowns are u_j and $u_{j^{\mathsf{c}}}$, so we may write

$$\underbrace{\begin{bmatrix} 1 & -U^\star_j (U^\star_{j^{\mathsf{c}}})^{-1} \end{bmatrix}}_{a^{\mathsf{T}}_\omega} \begin{bmatrix} u_j \\ u_{j^{\mathsf{c}}} \end{bmatrix} = 0,$$

or even as

$$au = 0,$$

where a has the entries of a_ω in the positions of ω and zeros elsewhere, i.e.

$$\underbrace{\begin{bmatrix} 1 & & -U_j^\star (U_{j^\mathsf{c}}^\star)^{-1} & \end{bmatrix}}_{\substack{a \\ 1 \times d}} \underbrace{\begin{bmatrix} \\ u_j \\ \\ u_{j^\mathsf{c}} \\ \\ \end{bmatrix}}_{\substack{u \\ d \times r}} = 0,$$

where the blank spaces represent zeros.

Notice that a is just as in §7.2, and as we mentioned there, a has exactly $r+1$ non-zero entries in the positions of ω. The 1 will always be in the j position, and $-U_j^\star (U_{j^\mathsf{c}}^\star)^{-1}$ in the j^c positions. Also, a depends on $S^\star$, but the choice of the basis $U^\star$ will only —if anything— scale a. In other words, $au = 0$ will always describe the same system of equations, no matter the choice of $U^\star$ and j.

In conclusion, since j in our discussion was an arbitrary element of ω, we conclude that ω constrains one entry of u_ω, namely u_j, as a function of $S^\star$ and the other entries of u_ω, namely u_{j^c}, through the equation $au = 0$. By constraining such entry, we are making sure that u_ω is *aligned* with $S^\star_\omega$, i.e., we are making sure that u_ω fits in $S^\star_\omega$.

Example 19. With $r = 2$,

$$\omega = \begin{bmatrix} \times \\ \times \\ \cdot \\ \times \\ \cdot \\ \cdot \\ \cdot \end{bmatrix} \quad \Longrightarrow \quad u = \begin{bmatrix} u_{j^\mathsf{c}} \\ \hline \\ \hline u_j \\ \hline \\ \\ \end{bmatrix} \longleftarrow \text{Determined given } u_{j^\mathsf{c}} \text{ and } S^\star.$$

■

Under the setup of §7.3, each linearly independent a_i is giving us an linearly independent vector orthogonal to a distinct projection of $S_i^\star$ onto a distinct $\mathbb{R}^d_{\omega_i}$. Recall that $\hat{S}_i$ is the projection of S onto $\mathbb{R}^d_{\omega_i}$. So when we constrain the vectors $u \in \mathcal{U}_\Omega$ through multiple a_i's, essentially what we are doing is making sure that the projections of $\mathcal{U}_\Omega$ and $S_i^\star$ onto $\mathbb{R}^d_{\omega_i}$ are the same *simultaneously* for every i.

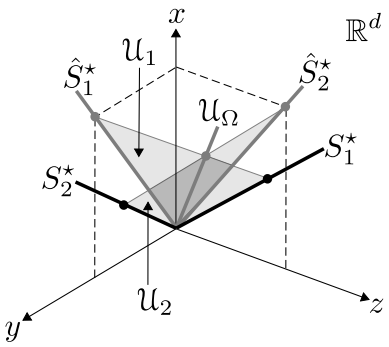

Figure 10: With the same setup as in Example 11, ω_1 would project onto the (x, y)-plane and ω_2 onto the (x, z)-plane. So when we use A to constrain several entries of u, basically what we are doing is making sure that the projection of $\mathcal{U}_\Omega$ and $S_1^\star$ onto $\mathbb{R}^d_{\omega_1}$, and the projections of $\mathcal{U}_\Omega$ and $S_2^\star$ onto $\mathbb{R}^d_{\omega_2}$ are the same.

Under the setup of this section, every $u \in \mathcal{U}_\Omega$ must satisfy $u_{\omega_i} \in S^\star_{\omega_i}$, for every $i \in \mathcal{I}$. Since each ω_i constrains one entry of u_ω, any $u \in \mathcal{U}_\Omega$ will now have several constrained entries, determined as functions of other r entries of u through the equation $Au = 0$, with A as defined in §7.3. How many constrained entries will $u \in \mathcal{U}_\Omega$ have? As many as independent ω's are contained in Ω, knowing that there are at most $d - r$.

In other words, each independent ω_i is constraining one entry of u_{ω_i}, say u_{j_i}, according to the other r entries of u_{ω_i}, say $u_{j_i^\mathsf{c}}$. And no matter how many ω_i's we have, there can be at most $d-r$ independent ones, i.e., we cannot constrain more than $d - r$ u_j's; this is formalized in Corollary 3.

Example 20. With $r = 2$,

$$\{\omega_1, \omega_2\} = \begin{bmatrix} \times & \cdot \\ \times & \cdot \\ \times & \cdot \\ \cdot & \cdot \\ \cdot & \times \\ \cdot & \times \\ \cdot & \times \end{bmatrix} \implies u = \left[\begin{array}{c} u_{j_1^\mathsf{c}} \\ \hline u_{j_1} \\ \hline \\ \hline u_{j_2^\mathsf{c}} \\ \hline u_{j_2} \end{array}\right] \begin{array}{l} \\ \longleftarrow j_1.\ \text{Determined by } S^\star_{\omega_1} \text{ and } u_{j_1^\mathsf{c}}. \\ \\ \\ \longleftarrow j_2.\ \text{Determined by } S^\star_{\omega_2} \text{ and } u_{j_2^\mathsf{c}}. \end{array}$$

■

9.1 Using Lemma 6

Intuitively, what §7.4 is telling us is that every independent ω constrains an entry of every $u \in \mathcal{U}_\Omega$ out of the $d - r$ that may be constrained, while §7.5 says that iff we can constrain $d - r$ entries of u, we will be guaranteed that there is only one r-dimensional subspace that fits Ω.

In other words, all we need to do is find an independent set of size $d - r$. Lemma 6 then completes the picture by telling us that a set $\breve{\Omega}$ is independent iff $m \geq n + r$ for every $\Omega \subset \breve{\Omega}$.

Example 21. Consider the same setup as in Example 11, we can see that $\mathbf{\Omega}$ is the only possible set with $d - r$ ω's. One can trivially verify that $m \geq n + r$ for every $\Omega \subset \mathbf{\Omega}$, hence there is only one r-dimensional subspace that fits $\mathbf{\Omega}$.

Intuitively, what this means is that given whichever entry of $u \in \mathcal{U}_\Omega$, one of the remaining entries of u is constrained according to $S^\star_{\omega_1}$ and an other of the remaining entries of u is constrained according to $S^\star_{\omega_2}$.

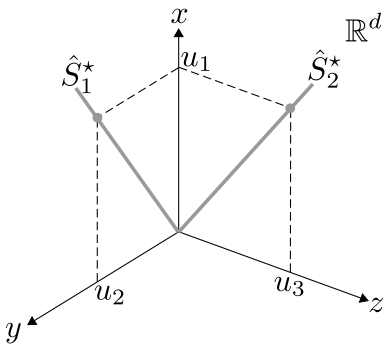

Figure 11: Given any entry of u, its remaining ones are constrained according to $S^\star_{\omega_1}$ and $S^\star_{\omega_2}$ in order to guarantee that u_{ω_i} fits in $S^\star_{\omega_i}$ for every i.

For example, if u_1 is given, u_2 and u_3 are constrained according to $S^\star_{\omega_1}$ and $S^\star_{\omega_2}$.

$$U = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix} \begin{array}{l} \longleftarrow \text{Given.} \\ \longleftarrow \text{Constrained according to } S^\star_{\omega_1}. \\ \longleftarrow \text{Constrained according to } S^\star_{\omega_2}. \end{array}$$

Explicitly, using (9) we can obtain:

$$\begin{aligned} u_2 &= U^\star_2 u_1 / U^\star_1, \\ u_3 &= U^\star_3 u_1 / U^\star_1. \end{aligned}$$

■

Conversely, $\mathbf{\Omega}$ is only constraining fewer than $d - r$ entries of u if there is no such Ω, i.e., there would be least one *free* entry of u to choose arbitrarily, whence there would be infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$.

Example 22. Consider $r = 2$ and

$$\mathbf{\Omega} = \begin{bmatrix} \times & \cdot & \times & \cdot \\ \times & \times & \cdot & \cdot \\ \times & \times & \times & \cdot \\ \cdot & \times & \times & \times \\ \cdot & \cdot & \cdot & \times \\ \cdot & \cdot & \cdot & \times \end{bmatrix}.$$

Again, $\mathbf{\Omega}$ is the only possible set with $d-r$ ω's. Nevertheless, $\mathbf{\Omega}$ does not satisfy the condition that $m \geq n+r$ for every one of its subsets Ω. Specifically, we can see that $\Omega = \{\omega_1, \omega_2, \omega_3\}$ is a subset of $\mathbf{\Omega}$ with $m < n+r$. By Theorem 2, there are infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$.

The fact that $\Omega = \{\omega_1, \omega_2, \omega_3\}$ satisfies $m < n+r$ implies that one of the ω's of Ω is redundant. This implies that there are at most $d-r-1$ independent ω's in $\mathbf{\Omega}$. Hence, there will be at least one *free* entry of u to choose arbitrarily —for example, one of the last two— and so there are infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$. ■

9.2 The idea behind Lemma 6

The condition of Lemma 6 is extremely simple and concrete; it depends only on the most elemental invariants of Ω: essentially, cardinalities of its subsets. In this section we will explain the intuition that led to this lemma.

One of the directions, that $\breve{\Omega}$ is dependent if $m < n+r$ for some $\Omega \subset \mathbf{\Omega}$, is directly given by Corollary 3. The idea behind it is that if we think of n as the number of equations in our system $Au = 0$, and m as the number of unknowns, since $\dim \ker A \geq r$ (otherwise there would be no r-dimensional subspace that fits Ω, as every r-dimensional subspace that fits Ω is contained in $\mathcal{U}_\Omega = \ker A$), then A must have at most $m-r$ linearly independent equations. If it has more than $m-r$ equations, then it must contain some linearly dependent ones, or else $\dim \ker A < r$.

The other direction, that if $\breve{\Omega}$ is dependent, it must contain a subset Ω with $m < n+r$, is essentially given by Lemma 7. We will spend the rest of this section explaining the intuition behind the later.

On one hand, we have that if Ω is a basis of ω, then ω is minimally dependent on Ω. On the other one, we know that ω determines u_j as a function of u_{j^c} through $au = 0$, so ω being dependent on Ω intuitively means that Ω also determines u_j as a function of u_{j^c} and $\mathcal{S}^\star$ through $Au = 0$.

The intuition behind Lemma 7 is that given u_{j^c}, if Ω is a basis of ω, u_j cannot be determined by Ω if there is at least one undetermined entry of u in the positions of $\mathcal{J} \backslash j^\mathsf{c}$. Why would this be true? Well, let u_{j^c} be given. If Ω is a trivial basis, it is clear that j is the only element of $\mathcal{J} \backslash j^\mathsf{c}$, and u_j is therefore determined by Ω given u_{j^c}.

If Ω is not a trivial basis, since Ω is independent, Ω constrains n entries of $u_{\mathcal{J} \backslash j^\mathsf{c}}$. Let $\boldsymbol{J}$ denote the positions of such entries. Explicitly, each $\omega_i \in \Omega$

determines one entry of $u_{\boldsymbol{J}}$, say u_{j_i}. Fix this correspondence, and let $j_i^{\mathsf{c}} := \omega_i \backslash j_i$.

Since ω is dependent on Ω, $j \in \boldsymbol{J}$, i.e., there must be some $\omega_{I_0} \in \Omega$, that determines u_j as a function of $u_{j_{I_0}^{\mathsf{c}}}$. But how could ω_{I_0} determine u_j if there is an undetermined entry in $u_{j_{\mathcal{I}_0}^{\mathsf{c}}}$? It can't, unless some of the $u_{j_i^{\mathsf{c}}}$'s cancel out, which almost surely won't happen. Then almost surely $u_{j_{I_0}^{\mathsf{c}}}$ must be determined, i.e., $j_{I_0}^{\mathsf{c}} \subset \boldsymbol{J}$.

Now it is convenient to define $J_0 := j$, and for $t > 0$, I_t as the set of i's such that ω_i determines an entry of u in the positions of J_t, and J_t as the set of positions of u that have not yet been determined at time $t-1$ and belong to some of the ω_i with $i \in I_{t-1}$. More precisely,

$$J_t := \left(\bigcup_{i \in I_{t-1}} \omega_i \right) \setminus \left(\bigcup_{\tau=0}^{t-1} J_\tau \right) \setminus j^{\mathsf{c}},$$
$$I_t := \{ i : j_i \in J_t \cap \boldsymbol{J} \}.$$

To give some intuitive meaning to this, consider the Tanner graph defined by the set Ω with disjoint sets of *column* vertices $\mathcal{I}$ and *row* vertices $\mathcal{J}$, where there is an edge between row vertex $j \in \mathcal{J}$ and column vertex $i \in \mathcal{I}$ if $j \in \omega_i$, and there is a special correspondence between $i \in \mathcal{I}$ and $j_i \in \boldsymbol{J}$ indicating that ω_i determines u_{j_i}.

For example, one such graph could look like the following, where bold edges represent the correspondence between $i \in \mathcal{I}$ and $j_i \in \boldsymbol{J}$, dashed vertices represent j^{c}, and bold vertices represent the elements of $\boldsymbol{J}$.

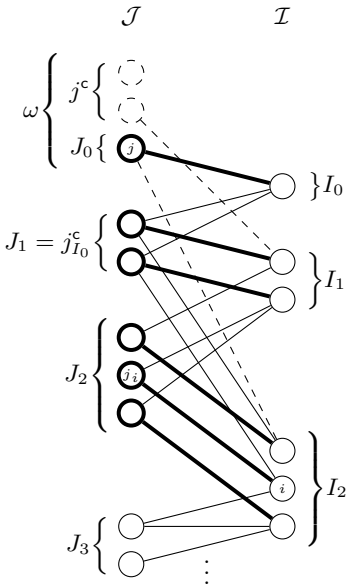

Since $J_1 = j_{I_0}^{\mathsf{c}}$, $J_t \subset \boldsymbol{J}$ for $t = 0, 1$. We can now do induction on t, and see that if there is an undetermined $j^* \in J_{t+1}$, there will almost surely be an

undetermined $j_i \in J_t$, as there will be an $i \in I_t$ determining u_{j_i} as a function of $u_{j_i^{\mathsf{c}}}$. And how can u_{j_i} be determined, if it is a function of $u_{j_i^{\mathsf{c}}}$, an there is an undetermined entry in $u_{j_i^{\mathsf{c}}}$, namely j^*? Again, it can't, unless some of the $u_{j_i^{\mathsf{c}}}$'s cancel out, which almost surely won't happen. Then almost surely $u_{J_{t+1}}$ must be determined, i.e., $J_{t+1} \subset \boldsymbol{J}$.

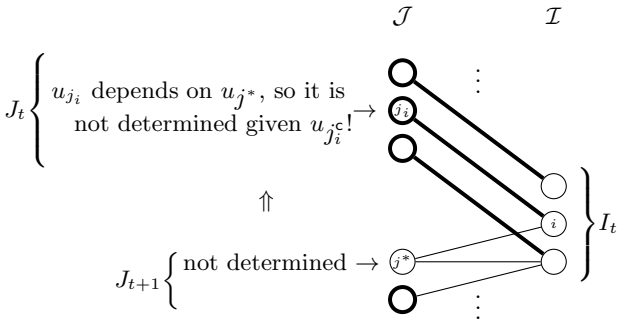

We will see that $\mathcal{J}\backslash j^{\mathsf{c}}$ is contained in the union of the J_t's, and since all the entries of u in the positions of the J_t's are determined, we will conclude that all the entries of $u_{\mathcal{J}\backslash j^{\mathsf{c}}}$ are determined. Since it takes one ω to constrain one entry of u, Ω must have $|\mathcal{J}\backslash j^{\mathsf{c}}| = m - r$ sets, i.e., $n = m - r$, which is precisely the statement of the lemma. To see this, let $\mathcal{J}_t$ and $\mathcal{I}_t$ denote the unions of the J_t's and the I_t's, respectively.

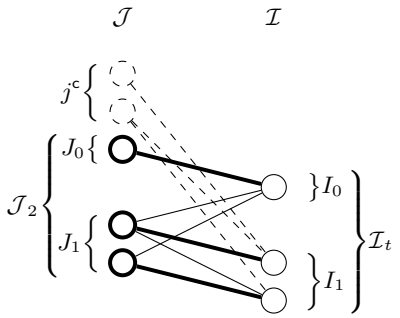

Let $\breve{\Omega} = \{\omega_i : i \in \mathcal{I}_t\}$, and notice that u_j is determined by $\breve{\Omega}$, which implies ω is dependent on $\breve{\Omega}$. Moreover, since $\mathcal{I}_t \subset \mathcal{I}$, $\breve{\Omega} \subset \Omega$. Since ω is minimally dependent on Ω, this means that $\breve{\Omega} = \Omega$, i.e., $\mathcal{I}_t = \mathcal{I}$, hence $\mathcal{J} = \mathcal{J}_t \cup j^{\mathsf{c}}$. Notice that $|J_t| = |I_t|$ for every t, so $|\mathcal{J}_t| = |\mathcal{I}_t|$. This implies

$$m = |\mathcal{J}| = \underbrace{|\mathcal{J}_t|}_{|\mathcal{I}_t|} + |j^{\mathsf{c}}| = \underbrace{|\mathcal{I}_t|}_{|\mathcal{I}|} + r = n + r,$$

as desired.

Lemma 6 then comes as a direct consequence of Lemma 7 by noticing that if $\boldsymbol{\Omega}$ contains a dependent ω, it must contain a basis Ω, and then the union of ω and Ω will satisfy $m < n + r$.

9.3 About our assumptions

It is clear that **A2**-**A4** are our only *real* assumptions, if any. And despite them being so lenient, for completeness we give a further discussion about them in this section. Specially to highlight how elemental they are, describe their tight relation, and talk about their generalizations.

We start with **A4**. Let us not loose sight of our final goal: whenever S fits $\hat{\Xi}$, we want to use $\hat{\mathfrak{X}}$ to validate if $\hat{\Xi}$ really lies in S, just as we would use the complete χ (see §2): by saying that if S also fits $\hat{\mathfrak{X}}$, it is because both, $\hat{\mathfrak{X}}$ and $\hat{\Xi}$ really lie in S.

If $|\omega| \leq r$, by our non-degenerate assumptions **A3** and **A6**, $S_\omega = \mathbb{R}^{|\omega|} = S^\star_\omega$ for every $S^\star \in \mathcal{S}^\star$. This implies that every $\hat{x}$ will fit in S, whether or not x or $\Xi_{\hat{x}}$ really lie in S, which totally defeats the purpose of $\hat{x}$. In other words, $\hat{x}$ would be completely useless.

Example 23. With $U_1^\star$ and $U_2^\star$ as in Example 11, if $|\omega| \leq r$, i.e., if $\omega = \{j\}$ or $\omega = \emptyset$, it is easy to see that any non-degenerate S —in this case any subspace spanned by a U with no zero entries— would fit every $\hat{x}$, no matter which subspace it belongs to. ■

Of course, having such an ω would not harm us at all; we could simply ignore it. This way, **A4** is a mere formality stating without loss of generality that we do not have useless information, knowing, of course, that if we had some, being so trivial to identify an ω with $|\omega| \leq r$, we could simply discard it.

On the other hand, if $|\omega| > r + 1$ it can only be easier to determine the missing entries of $\hat{x}$, and to determine if x really belongs to S, as any subspace that fits a larger ω will have to satisfy more restrictions. More precisely, there are fewer vectors that satisfy $u_{\omega_1} \in S^\star_{\omega_1}$ than vectors that satisfy $u_{\omega_2} \in S^\star_{\omega_2}$ if $|\omega_1| > |\omega_2|$, whence fewer vectors in $\mathcal{U}_1$ than in $\mathcal{U}_2$. In other words, there can only be fewer subspaces that fit $\hat{x}_1$ than subspaces that fit $\hat{x}_2$ if $|\omega_1| > |\omega_2|$.

And how could we use this information? Well, just as with χ, for a generic $x_\omega \in S^\star_\omega$ with $|\omega| > r$, we have that

(a) There is only one r-dimensional subspace of $\mathbb{R}^{|\omega|}$ that fits x_ω.

(b) Obvious, but essential, this single x lies in one and only one of the subspaces of $\mathcal{S}^\star$

Using this observation, we can easily generalize **A4** as follows: if $|\omega| > r + 1$, we could *split* ω into Ω, a set of $|\omega| - r + 1$ sets, each of size $r + 1$, that satisfies (1), such that Ω *behaves* just as ω in the sense that if there is an r-dimensional subspace that fits Ω:

(a) There is only one r-dimensional subspace of $\mathbb{R}^{|\omega|}$ that fits Ω, just as there is only one r-dimensional subspace of $\mathbb{R}^{|\omega|}$ that fits ω.

(b) $k_i = k_{\ddot{\imath}}$ for every $i, \ddot{\imath} \in \mathcal{I}$, just as there is only one k for ω.

Example 24. Suppose $r = 2$. Then we can split $\omega_1, \omega_2 \in \mathbf{\Omega}$ below as:

$$\mathbf{\Omega} = \begin{bmatrix} \times & \cdot \\ \times & \cdot \\ \times & \cdot \\ \times & \times \\ \times & \times \\ \cdot & \times \\ \cdot & \times \\ \cdot & \times \end{bmatrix} \sim \begin{bmatrix} \times & \cdot & \cdot & \times & \cdot & \cdot & \cdot & \\ \times & \times & \cdot & \cdot & \cdot & \cdot & \cdot & \\ \times & \times & \times & \cdot & \cdot & \cdot & \cdot & \\ \cdot & \times & \times & \times & \times & \cdot & \cdot & \times \\ \cdot & \cdot & \times & \times & \times & \times & \cdot & \cdot \\ \cdot & \cdot & \cdot & & \times & \times & \times & \cdot \\ \cdot & \cdot & \cdot & & \cdot & \times & \times & \times \\ \cdot & \cdot & \cdot & & \cdot & \cdot & \times & \times \end{bmatrix}.$$

Ω_1 Ω_2

Notice that both, Ω_1 and Ω_2 satisfy the conditions of Theorem 1. ■

We would thus obtain several orthogonal directions from ω rather than just one; more precisely, each ω_i would define a $(|\omega_i| - r + 1) \times d$ matrix, say A_i, where each subset of $|\omega_i| - r$ rows of A_i is full-rank. We could then use A_i just as a_i in §7.3, and construct A as

$$A = \begin{bmatrix} A_1 \\ \vdots \\ A_N \end{bmatrix},$$

and all the statements would work as before, with some straightforward adaptations.

This way, rather than an assumption, $|\omega| > r$ in **A4** is a requirement, while $|\omega_i| = r + 1 \ \forall \ i$ is a convenience that simplifies our notation, but at the same time forces us to work with the *minimal* possible assumption on $|\omega|$.

Now let us focus on **A3**. In §7.2 we mentioned (Lemma 3) that the orthogonal direction of $S^\star_\omega$, a_ω, has exactly $r + 1$ entries, which we just confirmed in §9. But why do we even care that a_ω has all non-zero entries? And what about the subspaces whose orthogonal direction has at least one non-zero entry? The answers to these questions are tightly related to the concept of degenerate subspaces, and help to give a better understanding of **A3** and **A4**.

An other interpretation of Lemma 3 is that all hyperplanes that have an orthogonal direction a_ω with at least one zero —the hyperplanes *aligned* with the *canonical* hyperplanes— are degenerate.

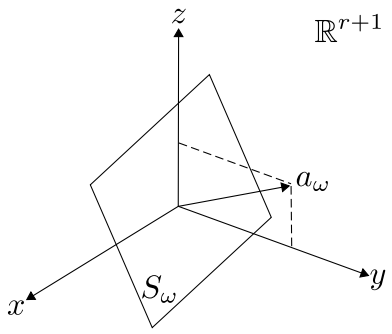

Figure 12: The planes of $\mathbb{R}^3$ that have an a_ω with at least one zero —the ones *aligned* with the *canonical* planes, (x, y) (x, z) or (y, z)— are degenerate.

So, if a_ω had at least one zero entry, $S^\star$ would have to be degenerate. And how would that affect us? Well, if $S^\star$ were degenerate, and $U^\star_{j^\mathsf{c}}$ turned out to be rank-defficient, $(U^\star_{j^\mathsf{c}})^{-1}$ would not exist, and we wouldn't be able to represent u_j as a function of u_{j^c} and $S^\star$ (see §9).

Example 25. For an example of a degenerate subspace, suppose $S^\star_\omega$ is the (y, z) plane in $\mathbb{R}^3$ and

$$u = \begin{bmatrix} u_1 \\ u_2 \\ u_3 \end{bmatrix}, \quad \omega = \begin{bmatrix} \times \\ \times \\ \cdot \end{bmatrix}, \quad \text{such that} \quad \hat{u} = \begin{bmatrix} 0 \\ u_2 \\ \cdot \end{bmatrix}.$$

Then, given $\hat{u}$, one could say nothing about u_3, as it could potentially be any point in the (y, z) plane.

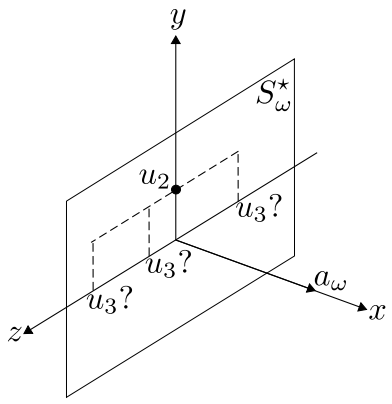

Figure 13: Given $\hat{u}$, one could say nothing about u_3, as it could potentially be any point in the (y, z) plane.

■

A little thought shows that under **A3**: there are no subspaces of dimension lower than r that fit ω; there are always infinitely many subspaces of dimension larger than r that fit ω; there are infinitely many subspaces of dimension larger than r that fit Ω if there is at least one r-dimensional subspace that fits Ω, and there is always a trivial subspace that fits Ω: $\mathbb{R}^d$.

Generalizing **A2** becomes straightforward with these observations: if dimensions of the subspaces in $\mathcal{S}^\star$ are unknown simply, iterate over $r = 1, ...d$. But more importantly, these observations tell us what would happen if $S^\star$ were degenerate, i.e., if we dropped **A3**. First, there would be infinitely many r-dimensional subspaces that fit $\mathbf{\Omega}$, no matter what $\mathbf{\Omega}$ is; even if $\{1, ..., d\} \in \mathbf{\Omega}$, just as there could be infinitely many r-dimensional subspaces that fit a generic χ from an $(r-1)$-dimensional subspace.

But a degenerate r-dimensional subspace is just a subspace that is even lower-rank *in some sections*; the *dimensions* corresponding to $\upsilon \subset \{1, ..., d\}$ with $|\upsilon| \leq r$ such that $\dim S_\upsilon < |\upsilon|$. So if we wanted to determine $S^\star$ and whether or not our data really lies in S whenever S fits our data, we could also iterate over $\tau = 1, ..., r$ to find the lower-dimension *portion* of $S^\star$. It is clear that in such cases, an ω with $|\omega| \leq r$ would definitely be useful to identify the lower-dimensional *portions* of $S^\star$.

9.4 Almost every $\neq$ every

We close this section with an important reminder: our results hold for almost every $\mathcal{S}^\star$. This means that there are *bizarre* and *extremely unlikely* sets of subspaces for which our results do not hold.

To see what exactly we mean by *bizarre*, consider the converse of Theorem 1, stated precisely in Corollary 6. It states that if $\mathbf{\Omega}$ satisfies the conditions of the theorem, and $k_i \neq k_{\ddot{\imath}}$ for some $(i, \ddot{\imath})$, almost surely there is no possible way that the projections of some r-dimensional subspace S and $S^\star$ onto $\mathbb{R}^d_{\omega_i}$, namely $\hat{S}_i$ and $\hat{S}^\star_i$, are equal for every i. For this to happen, $\hat{S}^\star_i$ would have to be *aligned* with $\hat{S}^\star_{\ddot{\imath}}$ in the following sense.

Example 26. Suppose $k_1 \neq k_2 \neq k_3$, and $\mathbf{\Omega}$ is as in Example 18. Let's consider ω_1 and ω_2 first. The intersection of $\mathcal{U}_1$ and $\mathcal{U}_2$ is a line (see Figure 14).

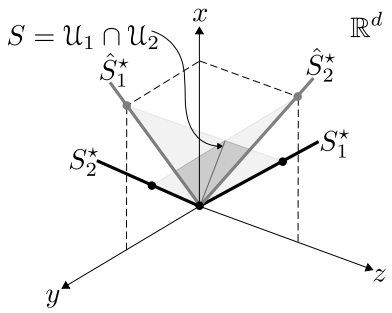

Figure 14: The intersection of $\mathcal{U}_1$ and $\mathcal{U}_2$.

Recall that there is an r-dimensional subspace that fits Ω only if $\dim \mathcal{U}_\Omega \geq r$, and $\mathcal{U}_\Omega = \bigcap_{i=1}^N \mathcal{U}_i$, so in this case $\dim \mathcal{U}_\Omega = r = 1$ only if $S_3^\star$ is *aligned* with $S_1^\star$ and $S_2^\star$ in the sense that the projections of $S_3^\star$ and $S = \mathcal{U}_1 \cap \mathcal{U}_2$ onto $\mathbb{R}^d_{\omega_3}$, in this case the (y, z)-plane, are the same. Of course, there is only a set of measure zero for which $S_3^\star$ will satisfy this, e.g., the set of subspaces in the shaded plane of Figure 15.

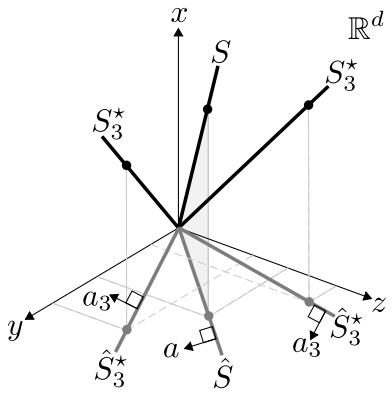

Figure 15: There is an r-dimensional subspace that fits Ω only if $S_3^\star$ lies in the shaded plane, which is a set of measure zero, as $S_3^\star \subset \mathbb{R}^3$; equivalently, if a_3 is aligned with a, which is also a set of measure zero, as $a_{\omega_3} \in \mathbb{R}^2$.

If $S_3^\star$ happened to satisfy this condition, we would have a situation as in Figure 16.

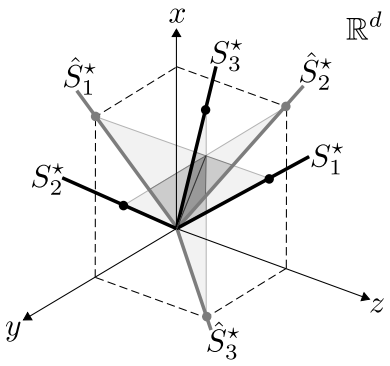

Figure 16: Set of measure zero in which $S_1^\star$, $S_2^\star$ and $S_3^\star$ are *aligned*.

In general, of course, the intersection of $d-r+1$ distinct hyperplanes in $\mathbb{R}^d$, is an $r-1$ subspace, as in Figure 17.

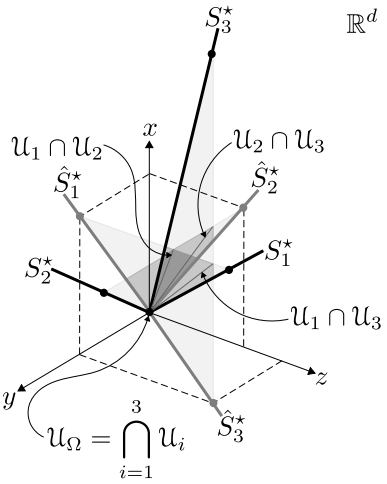

Figure 17: Almost surely, the intersection of $d-r+1$ distinct planes in $\mathbb{R}^d$, is an $(r-1)$-dimensional subspace. In our example, the intersection of the three different planes in $\mathbb{R}^3$ corresponding to each of the $\mathcal{U}_i$'s has dimension zero, i.e., it is just the origin.

■

A similar phenomenon occurs for Lemma 7, where Ω may be a basis for ω even if $m < n+r$. For this to happen, some projections of $S_i^\star$ would have to be *aligned* with some other projections of $S_i^\star$.

All this to summarize that our results hold for almost every $\mathcal{S}^\star$. For example, if $\mathcal{S}^\star$ were a collection of *random* subspaces, we would know that our results hold with probability 1.

10 Epilogue

We now have deterministic necessary and sufficient conditions to determine if certain incomplete data really *lies* in a subspace whenever it *fits* in a subspace. Moreover, we also have deterministic necessary and sufficient conditions to determine if there is only one subspace that *fits* certain partially observed data. We answer these questions by characterizing when and only when a set of incomplete vectors behaves as a single but complete one, in the sense described in §2.

Being such fundamental problems, it comes as no surprise that our results as well as the tools we used to derive them, have many powerful consequences and relevant applications; already have been discussed some of these in §3. Finally, it is presented a brief discussion of some other very interesting related problems for future research.

P1. This is one particularly interesting idea: could a set of incomplete vectors $\hat{\Xi}$ behave as a set of r complete ones, forming a basis for a subspace S, such that S could be uniquely determined only by $\hat{\Xi}$?

To be more precise, as we said in §3, in this document we assumed that a subspace S that fit certain incomplete data $\hat{\Xi}$ was given to us, and wanted to validate if $\hat{\Xi}$ really *lied* in S using $\hat{\mathcal{X}}$ as a validating vector. We now want to know what necessary and sufficient properties should $\hat{\Xi}$ have, e.g., what conditions on its observed entries, say Υ, such that S could be uniquely identified solely from $\hat{\Xi}$, in a deterministic way, just as it would be determined by a set of r complete vectors.

This is exactly one of the low-rank matrix completion open questions: what are the necessary and sufficient conditions on Υ to guarantee that the subspace where Ξ lies can be uniquely identified from $\hat{\Xi}$. We now know when and only when a set of generic incomplete vectors $\hat{\mathcal{X}}$ behaves as a complete generic vector in the sense described in §2. I think this might give some insight to finding necessary and sufficient conditions on Υ for deterministic low-rank matrix completion, and conjecture that

> *All that is really required is that $\hat{\Xi}$ contains a set of r sets of incomplete vectors $\hat{\mathcal{X}}_1, ..., \hat{\mathcal{X}}_r$ observed in $\boldsymbol{\Omega}_1, ..., \boldsymbol{\Omega}_r$ such that each $\boldsymbol{\Omega}_\tau$ satisfies the conditions of Theorem 1*

P2. It is also of interest to characterize precisely and in a very concrete and intuitive manner what are the characteristics of the set of measure zero under which our results do not hold.

P3. There is a relation between this work and concepts like the spark of a matrix, r-connected graphs, matroids and matchings, to mention some.

Using tools from the corresponding areas one may derive powerful corollaries of our results. For example, one corollary of Lemma 6 is that there

are infinitely many r-dimensional subspaces that fit $\boldsymbol{\Omega}$ if the Tanner graph defined by $\boldsymbol{\Omega}$, as described in §9.2, is not r-connected in the *row* vertices.

I would also like to pursue investigate these relations in more depth, with the hope of arriving at a non-combinatorial relaxation of the conditions of Lemma 6 that could be computationally efficient to verify.

P4. In some sense, in §2 we are reconstructing a subspace from its projections onto $\mathbb{R}^d_\omega$. I would also like to explore this idea more, and see if and when can we reconstruct a subspace from arbitrary projections.

P5. Observe that $\boldsymbol{A}u = 0$ is essentially describing a linear variety defined by N equations and d unknowns, where each equation only involves a subset of the unknowns; $\boldsymbol{\Omega}$ is the set describing which variables are involved in which equation, i.e., the *support* of equation i are the variables u_{ω_i}.

Theorem 2 bounds the dimension of such variety only as a function of the supports of its equations; I believe that it can also be used to bound the dimension of similar generic algebraic varieties, not necessarily linear, based only on the support of each of the polynomials that define it.

11 List of symbols, references and index

Symbol	Description	pp.
$\boldsymbol{A}$	$N \times d$ matrix with $\{a_i\}_{i=1}^{N}$ as its rows	25
A	$n \times d$ matrix with $\{a_i\}_{i \in \mathcal{I}}$ as its rows.	25
a_ω	Vector in $\mathbb{R}^{r+1}$ orthogonal to $S^\star_\omega$.	25
a	$1 \times d$ row vector with the entries of a_ω in the positions of ω.	25
d	Ambient dimension.	11
$\mathcal{I}$	Indices of the ω's in $\boldsymbol{\Omega}$ that belong to Ω.	12
i, $\ddot{\imath}$	Used to index vectors. In general, $\in \{1, ..., N\}$.	11
$\mathcal{J}$	Observed *rows* of Ω.	12
j	Used to index entries or elements. In general, $\in \{1, ..., d\}$.	12
j^{c}	Given j and ω, $j^{\mathsf{c}} := \omega \backslash j$.	40
$\mathcal{K}$	Multiset of indices $k_i \in \{1, ..., K\}$ indicating that $x_i \in S^\star_{k_i}$.	11
K	Number of subspaces in $\mathcal{S}^\star$.	11
k	Used to index *distinct* subspaces. In general, $\in \{1, ..., K\}$.	11
k_i	Index denoting which subspace of $\mathcal{S}^\star$ x_i belongs to.	11
ℓ	Number of linearly independent rows in A.	??
m	Number of distinct observed *rows* in Ω.	12
N	Number of vectors in $\mathcal{X}$.	11
n	Number of sets that Ω contains.	12
$\boldsymbol{\Omega}$	Set of observed entries of $\hat{\mathcal{X}}$.	12
Ω	Subset of $\boldsymbol{\Omega}$.	12
ω	Set of observed entries of $\hat{x}$. Subset of $\{1, ..., d\}$ of size $r+1$	12
r	Dimension of the subspaces.	11
$\mathbb{R}^d_\omega$	Subspace of $\mathbb{R}^d$ spanned by the canonical vectors. corresponding to the positions of ω	13
S	Arbitrary subspace, generally r-dimensional.	11
$\hat{S}$	Projection of S onto $\mathbb{R}^d_\omega$.	13
S_ω	Subspace of $\mathbb{R}^{\lvert\omega\rvert}$. The restriction of S to ω.	13
$S^\star$	Arbitrary subspace of $\mathcal{S}^\star$.	11
$S^\star_k$	k^{th} subspace of $\mathcal{S}^\star$.	11
$S^\star_i$	Subspace where x_i lies. ${k_i}^{th}$ subspace of $\mathcal{S}^\star$.	11
$\hat{S}^\star$	Projection of $S^\star$ onto $\mathbb{R}^d_\omega$.	13
$S^\star_\omega$	Subspace of $\mathbb{R}^d$. The restriction of $S^\star$ to ω.	13
$\mathcal{S}^\star$	Set of K r-dimensional subspaces. $\mathcal{X}$ lies in their union.	11
$\mathcal{U}_\Omega$	Subspace of all r-dimensional subspaces that fit Ω.	23
$\mathcal{U}_\omega$	Subspace of all r-dimensional subspaces that fit ω.	23
$\mathcal{U}_i$	Subspace of all r-dimensional subspaces that fit ω_i.	25
U	Arbitrary basis of S.	11
$\hat{U}$	Arbitrary basis of $\hat{S}$.	13
U_ω	Arbitrary basis of S_ω.	13
$U^\star$	Arbitrary basis of $S^\star$.	11

Continued on next page

Table 1 – *Continued from previous page*

Symbol	Description	pp.
$\hat{U}^\star$	Arbitrary basis of $\hat{S}^\star$.	13
$U^\star_\omega$	Arbitrary basis of $S^\star_\omega$.	13
u	Arbitrary vector, typically from S, $\mathcal{U}_\omega$ or $\mathcal{U}_\Omega$.	23
$\hat{u}$	Incomplete version of u; analogous to $\hat{x}$.	13
u_ω	Vector in $\mathbb{R}^{\lvert\omega\rvert}$ with the entries of u in the positions of ω.	13
$\mathcal{X}$	Set of N vectors in the union of the subspaces in $\mathcal{S}^\star$.	11
$\hat{\mathcal{X}}$	Incomplete version of $\mathcal{X}$.	13
$\hat{\Xi}$	Set of incomplete vectors that fit in an r-dimensional subspace.	5
x	Arbitrary vector of $\mathcal{X}$.	11
x_i	i^{th} vector of $\mathcal{X}$.	11
$\hat{x}$	Incomplete version of x.	13
x_ω	Vector in $\mathbb{R}^{\lvert\omega\rvert}$ with the entries of x in the positions of ω.	13
χ	Complete generic vector.	5
$\times$	Symbol to denote that an entry is observed.	12
$\cdot$	Symbol to denote that an entry is missing.	12
$\breve{\circ}$	Analogous to $\circ$, e.g., $\breve{n} = \breve{\Omega}$	12
$\bar{\circ}$	Analogous to $\circ$, e.g., $\bar{m}$ is the number of observed *rows* of $\breve{\Omega}$	12
$\circ_{\hat{i}}$	Analogous to $\circ$, e.g., $\hat{S}^\star_i$ is the projection of $S^\star_i$ onto $\mathbb{R}^d_{\omega_i}$	13

References

[1] René Vidal, "Subspace clustering," *IEEE Signal Processing Magazine*, 2010.

[2] Benjamin Recht, "A simpler approach to matrix completion," *Jrnl. of Machine Learning Rsrch.*, 2011.

[3] Laura Balzano, Bejamin Recht, and Robert Nowak, "High-Dimensional Matched Subspace Detection When Data are Missing," in *Proceedings of the International Conference on Information Theory*, June 2010.

[4] Daniel Pimentel, Laura Balzano, and Robert Nowak, "On the sample complexity of subspace clustering with missing data," in *Proceedings of the Statistical Signal Processing Workshop*, 2014.

[5] Brian Eriksson, Laura Balzano, and Robert Nowak, "High-Rank Matrix Completion and Subspace Clustering with Missing Data," in *Proceedings of the Conference on Artificial Intelligence and Statistics (AI Stats)*, 2012.

[6] Stephen H. Friedberg, Arnold J. Insel, and Lawrence E. Spence, "Linear algebra," 2003.

[7] James Oxley, "Matroid theory," 2011.

[8] W. Murray Wonham, "Linear multivariable control," 1979.

Index

$\mathbb{R}^d_\omega$, 13

basis, 28
behave, 5

degenerate subspace, 14
dependent set Ω, 27
dependent set ω, 27

incomplete vector $\hat{x}$, 13
independent set Ω, 27
independent set ω, 27

projected basis $\hat{U}$, 13
projected subspace $\hat{S}$, 13

redundant set, 27
restricted basis U_ω, 13
restricted subspace S_ω, 13
restricted vector x_ω, 13

to fit a generic vector, 15
to fit Ω, 16
to fit ω, 15
to fit $\hat{\mathfrak{X}}$, 15
to fit $\hat{x}$, 15